arXiv:1703.03478v1 [cs.LG] 9 Mar 2017

On-line Learning with Abstention

Corinna Cortes [1] **Giulia DeSalvo** [2] **Mehryar Mohri** [2 1] **Scott Yang** [2]

Abstract

We introduce and analyze an on-line learning setting where the learner has the added option of abstaining from making a prediction at the price of a fixed cost. When the learner abstains, no feedback is provided, and she does not receive the label associated with the example. We design several algorithms and derive regret guarantees in both the adversarial and stochastic loss setting. In the process, we derive a new bound for on-line learning with feedback graphs that generalizes and extends existing work. We also design a new algorithm for on-line learning with sleeping experts that takes advantage of time-varying feedback graphs. We present natural extensions of existing algorithms as a baseline, and we then design more sophisticated algorithms that explicitly exploit the structure of our problem. We empirically validate the improvement of these more sophisticated algorithms on several datasets.

1. Introduction

We consider an on-line learning scenario, crucial in many applications, where the learner is granted the option of abstaining from making a prediction. In classification for example, at each round, the learner can choose to make a prediction and incur a fixed misclassification cost in the case of an inaccurate prediction, or elect to abstain, in which case she incurs a fixed abstention cost, typically lower than the misclassification cost. Abstention can thus represent an attractive option to avoid a higher cost of misclassification. Note, however, that in the case of abstention, the learner is not given the true label (correct class), which therefore results in an information loss.

The scenario of on-line learning with abstention appears in many important applications. As an example, consider the medical diagnosis setting. A doctor can choose to make a diagnosis based on the current information available about a patient, or abstain and request further laboratory tests, which can represent both a time delay and a financial cost. The abstention cost of course is typically substantially lower than that of a wrong diagnosis. The on-line model is appropriate since it captures the gradual experience a doctor gains by examining and following new patients. The adversarial setting may correspond to highly specialized medical facilities that receive a small and difficult subset of patients that cannot be naturally described by a stochastic model. The stochastic setting of on-line learning with abstention could model the case of general physicians who see patients from a wide pool containing both easy and hard to classify diseases.

Another instance of this problem appears in the design of spoken-dialog applications such as modern personal assistants. Each time the user asks a question, the assistant can either provide a direct response to the question or choose to say "I am sorry, I do not understand", in which case the user must either reformulate his question or pursue a longer dialog. This reaction is costly but typically less than that of receiving a wrong piece of information. Over time, the assistant adapts to the user's preferences and only abstains from very complex requests. Similar on-learning problems arise in the context of self-driving cars where, at each instant, the assistant must determine whether to continue steering the car or return the control to the driver.

On-line learning with abstention also naturally models many problems arising in electronic commerce platforms such as an Ad Exchange, an on-line system set up by a publisher where several advertisers bid in order to compete for an ad slot. An advertiser must first choose whether to participate in an auction for an ad slot and then determine the amount for the bid. The abstention cost in this application is the opportunity loss of not bidding for a specific ad slot. A similar problem arises in on-line advertisement.

In the batch setting, the problem of learning with abstention has been studied in several publications, including a consistency analysis and a related algorithm for confidence-based rejection by Bartlett & Wegkamp (2008). More recent comprehensive studies of the problem by Cortes et al. (2016a;b) include an analysis of the properties of the corresponding abstention (or rejection) loss with a series of theoretical guarantees and algorithmic results both for kernel-

[1]Google Research, New York, USA [2]Courant Institute, New York, USA. Correspondence to: Giulia DeSalvo <gd775@nyu.edu>.

based hypotheses and boosting stumps.

In this paper, we present an extensive study of on-line learning with abstention for both the adversarial and stochastic settings. We consider the common setting of prediction with expert advice (Littlestone & Warmuth, 1994), where we adopt the same loss function and hypotheses made up of predictor and abstention function pairs as in (Cortes et al., 2016a). However, a key aspect of our learning scenario that makes it distinct from both batch learning with abstentions and standard on-line learning problems is that, if at a given round, the algorithm abstains from making a prediction on the received input point, then the classification label of that point is not revealed. As a result, the loss of those experts who would not have opted to abstain from making a prediction is not known for that point. Thus, we are dealing with a *semi-bandit problem*: if the algorithm chooses to make a prediction, then the true label is revealed and the losses of all experts are known. But, if it elects to abstain, then only the losses of the abstaining experts are known, all of them equal to the fixed abstention loss.

Within the on-line learning literature, work related to our scenario includes the KWIK (*knows what it knows*) framework of Li et al. (2008) in which the learning algorithm is required to make only correct predictions but admits the option of abstaining from making a prediction. The objective is then to learn a concept exactly with the fewest number of abstentions. This can be seen as a special case of our framework for on-line learning with abstention with an infinite misclassification cost and some finite abstention cost. A relaxed version of the KWIK framework was introduced and analyzed by Sayedi et al. (2010) where a fixed number k of incorrect predictions are allowed with a learning algorithm related to the solution of the so-called "meta-egg game" puzzle. A theoretical analysis of learning in this framework was also recently given by Zhang & Chaudhuri (2016). Our framework does not strictly cover this relaxed framework. However, for some choices of the misclassification cost depending on the horizon, the framework is very close to ours. The analysis in these frameworks was given in terms of mistake bounds since the problem is assumed to be realizable. We will not restrict ourselves to realizable problems, and instead, will consider regret guarantees.

To devise learning algorithms for online learning with abstention, we draw from connections with several related on-line learning problems. As we shall see, our semi-bandit problem can be cast as a specific instance of on-learning with a feedback graph, a framework introduced by Mannor & Shamir (2011) and later further analyzed by several authors (e.g., Caron et al. (2012); Alon et al. (2013; 2014; 2015); Kocák et al. (2014); Neu (2015); Cohen et al. (2016)). In our context, the feedback graph varies over time, a scenario to which most of the existing algorithms or analyses do not readily apply. We also show that, remarkably, our problem can be connected to the sleeping experts setting introduced and analyzed by Kleinberg et al. (2008).

This paper is organized as follows. Section 2 presents a formal definition of our on-line learning problem for both stochastic and adversarial scenario, including the corresponding definitions of regret. In Section 3, we analyze the adversarial setting. We first show that on-learning with abstention can be cast as an on-line learning problem with a time-varying graph that we will describe in detail. We then leverage the algorithm of Alon et al. (2015) and the implicit exploration technique of Neu (2015) to design algorithms for on-line learning with abstention in this setting, for which we give regret guarantees that hold both in expectation and with high probability. Our bounds exploit the properties of the abstention loss and reflect the relative cost of abstention to misclassification specified in the scenario.

In Section 4, we analyze the stochastic setting of on-line learning with abstention. Stochastic bandits with a fixed feedback graph have been previous previously studied by Caron et al. (2012); Cohen et al. (2016). Note that the paper of Cohen et al. (2016) does not apply to our setting since they assume independence between graphs and losses. We first extend the work of Caron et al. (2012) and derive regret bounds that hold for time-varying feedback graphs, and then describe an application to on-line abstention. We further introduce a novel algorithm for this setting, UCB-D, that exploits the fact that some decisions (abstentions) incur a fixed cost: the abstention cost. Finally, we give an alternative algorithm for this setting, UCB-H, which makes use of the hierarchical structure of the abstention loss. We show that UCB-H can be viewed as an extension of a novel algorithm for sleeping experts with feedback graphs. Thus, we then derive new sleeping expert regret guarantees for time-varying feedback graphs.

We also report the results of multiple experiments with both artificial and real-world data sets comparing the performance of these algorithms (Section 5). We show that UCB-D and UCB-H both tend to outperform the baseline derived from Caron et al. (2012), and that UCB-H, in particular, is able to exploit the hierarchical nature of the problem to even achieve negative regret.

2. Learning problem

Let $\mathcal{X}$ denote the input space and let $\mathcal{E} = \{\xi_j = (h_j, r_j) \colon j \in [1, K]\} \subseteq \mathcal{H} \times \mathcal{R}$ denote a family of experts where r_j is an abstention function that determines which points to abstain on and h_j is a hypothesis function that makes the prediction. This type of joint family of hypothesis and abstention functions was first proposed in the batch setting by Cortes et al. (2016a). Then, consider

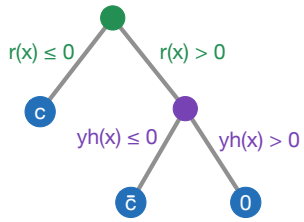

Figure 1: Hierarchical structure of the abstention loss.

the following sequential prediction problem. At each round $t \in [1, T]$, the learner receives an input $x_t \in \mathcal{X}$ and chooses (potentially at random) an index $I_t \in [K]$ corresponding to a hypothesis and abstention function pair, $\xi_{I_t} = (h_{I_t}, r_{I_t})$. The learner determines whether to make a prediction based on the abstention function, $r_{I_t}(x_t)$. One natural choice for $r(x)$ is a confidence-based function which is based on the magnitude of the hypothesis $h(x)$, that is $r(x) = |h(x)| - \gamma$ for some threshold, γ (Bartlett & Wegkamp, 2008). If $r_{I_t}(x_t) \leq 0$, the learner abstains and incurs a fixed loss $c \in \mathbb{R}_+$. If $r_{I_t}(x_t) > 0$, the learner accepts and makes a prediction $h_{I_t}(x_t)$. In this case, she receives a label $y_t \in \{\pm 1\}$ and incurs the prediction loss $\bar{c} 1_{y_t h_{I_t}(x_t) \leq 0}$, where $\bar{c} \in \mathbb{R}$. Let $z_t = (x_t, y_t)$. The instantaneous *abstention loss* of the learner is defined as

$$L(\xi_{I_t}, z_t) = \bar{c} 1_{y_t h_{I_t}(x_t) \leq 0} 1_{r_{I_t}(x_t) > 0} + c 1_{r_{I_t}(x_t) \leq 0}.$$

The input and label can either be drawn from an unknown distribution $\mathcal{D}$ or received adversarially. Common benchmarks in this setting are the expected regret

$$\mathrm{Reg}_T(\mathcal{A}) = \mathbb{E}\left[\max_{\xi \in \mathcal{E}} \sum_{t=1}^{T} L(\xi_{I_t}, z_t) - \sum_{t=1}^{T} L(\xi, z_t)\right]$$

and the pseudo-regret

$$\overline{\mathrm{Reg}}_T(\mathcal{A}) = \max_{\xi \in \mathcal{E}} \mathbb{E}\left[\sum_{t=1}^{T} L(\xi_{I_t}, z_t) - \sum_{t=1}^{T} L(\xi, z_t)\right].$$

The expectations are over the learner's actions I_t, and, in the stochastic setting, also over the choice of $z_t \sim \mathcal{D}$.

The abstention loss admits a hierarchical structure that can be characterized by a binary tree where the root node splits the inputs x_t according to the sign of r. More precisely, if $r(x) \leq 0$, then x is sent to the left child and the loss of x is c while if $r(x) > 0$, then x is sent to the right child, which is then split into two nodes according to the sign of $yh(x)$ (see Figure 1). We will see how two of our algorithms exploit this specific loss structure.

3. Adversarial scenario

Consider first online-learning with abstention in an adversarial scenario where we do not assume any distributional properties about the data (x_t, y_t).

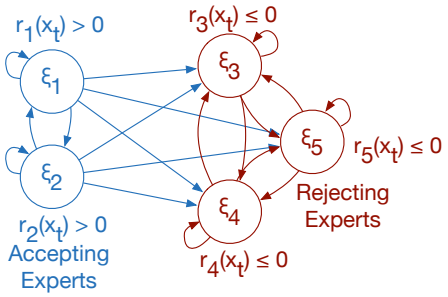

Figure 2: Feedback graph for on-line abstention.

We can interpret the abstention setting as an instance of on-line learning with feedback graphs. Recall that in the standard prediction with expert advice framework, the learner receives the loss of every expert at each round. On the other hand, in the multi-armed bandit framework, the learner receives only the loss of the chosen expert at each round. The on-line learning with feedback graph model is an intermediate scenario, where the loss observations associated to each expert are modeled using a directed graph.

Formally, let $G = (V, E)$ denote a directed graph, where E denotes the set of edges between the vertices V. Then, a directed time-varying feedback graph is a sequence of directed graphs $G_t = (V, E_t)$ over the set of experts $V = \{1, \ldots, K\}$ at round t such that if a vertex i admits a directed edge to vertex j, then the loss of vertex j is observed whenever vertex i is chosen. Moreover, recall that an independence set of a graph is a set of vertices $S_t \subseteq V$ that are not connected by any edges and the independence number $\alpha(G_t)$ is the size of the largest independence set.

The feedback graph of on-line learning with abstention contains a total of $|\mathcal{E}| = K$ vertices, one for each expert. If the learner accepts at round t, she observes the loss of all experts since the label is revealed. If she abstains at round t, she will only see the loss of the experts that are abstaining. Thus, the feedback graph is a nearly fully connected graph with self-loops, except that it has only one-way edges from the accepting nodes to the abstaining nodes. See Figure 2 for an illustration of the graph. Using this feedback graph, we can then adapt the EXP3.G algorithm from Alon et al. (2015) to the on-line abstention setting. This algorithm is similar to EXP3 except that its surrogate loss has an importance weighting scheme that is based on the probability of an expert being observed, as opposed to being chosen. For the pseudocode, see Algorithm 5 in Appendix A.1.

We now show an expected regret bound in terms of these feedback graphs and the abstention loss. Below, $K_{abs,t}$ denotes the number of vertices whose corresponding experts abstain on input x_t. Similarly $K_{acc,t}$ denotes the number of vertices whose experts accept on input x_t. Thus, $K_{acc,t} + K_{abs,t} = K$.

Theorem 1. *For any $T \geq 1$ and any $\xi^\star \in \mathcal{E}$,* EXP3.ABS *admits the guarantee:*

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)L_t(\xi)\right] - \sum_{t=1}^{T} L_t(\xi^\star) \leq \gamma T + \frac{\log K}{\eta} + \eta\sum_{t=1}^{T}\left(8c^2\log\frac{4KK_{abs,t}}{\gamma} + 8\bar{c}^2\log\frac{4KK_{acc,t}}{\gamma}\right).$$

The theorem shows a clear trade-off between the cost of abstention c and the cost of misclassification $\bar{c}$, scaled by the number of vertices that are abstaining and accepting respectively. By letting $\gamma = \min(\frac{1}{\sqrt{T}}, \frac{1}{2})$ and $\eta = 2\gamma$, the expected regret can be bounded by $O\big((c^2 + \bar{c}^2)\sqrt{T}\log(KT)\big)$. This regret bound in expectation can be upgraded to a high probability bound using the implicit exploration technique of Neu (2015), which applies a new surrogate loss with lower variance. The pseudocode of this algorithm, EXP3-IX.ABS, is given in Algorithm 6 in Appendix A.1.

Theorem 2. *For any $\delta > 0$, $\gamma \leq \frac{1}{\bar{c}K}$ and $\gamma' \leq \frac{1}{cK}$,* EXP3-IX.ABS *admits the following guarantee with probability at least $1 - 4\delta$:*

$$\begin{aligned}\sum_{t=1}^{T} L_t(\xi_{I_t}) - L_t(\xi^\star) &\leq \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} \\ &+ \left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right)\sum_{t=1}^{T} 2\bar{c}f(K_{acc,t}) \\ &+ \left(\frac{\eta}{2}c + \kappa + (e-2)\gamma'\right)\sum_{t=1}^{T} 2cf(K_{abs,t}) \\ &+ \frac{\log K}{\eta} + \frac{3\eta}{4}\frac{K\log\left(\frac{K}{\delta}\right)}{\kappa} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'},\end{aligned}$$

where $f(x) = \log\left(1 + \left\lceil\frac{x^2}{\kappa}\right\rceil + x\right) + 1$.

The cost of abstention and misclassification appear to affect the bound in this result in a similar way as in Theorem 1. The difference here is that instead of scaling these terms by $\bar{c}^2$ and c^2, the factors are more involved: $2\bar{c}\left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right)$ and $2c\left(\frac{\eta}{2}c + \kappa + (e-2)\gamma'\right)$. [Theorem 5] in Alon et al. (2014) provides tight lower bounds for the general setting of on-line learning with fixed feedback graphs. Since one can always design a specific case where the time-varying feedback graphs are fixed in the abstention scenario, the bounds presented above are also tight.

4. Stochastic scenario

In this section, we analyze the scenario where at each round, the input and label are drawn i.i.d. from some unknown distribution: $(x_t, y_t) \sim \mathcal{D}$. Caron et al. (2012) consider a setting of stochastic multi-armed bandits where the feedback is modeled by an undirected graph that is fixed over time. They design an algorithm, UCB-N, that runs the standard UCB algorithm, but updates the upper-confidence bound of every arm that receives feedback. The analysis in Caron et al. (2012) cannot be directly applied to the on-line learning with abstention setting, because the feedback graphs change over time.

ALGORITHM 1: UCB-N.ABS adapted from Caron et al. (2012)

for $t \geq 1$ **do**
  RECEIVE(x_t)
  $\xi_{I_t} \leftarrow \operatorname{argmin}_{\xi_j\in\mathcal{E}}\left\{\widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1}\right\}$
  if $r_{I_t}(x_t) > 0$ **then**
    REVEAL(y_t)
  for $\xi_j \in \mathcal{E}$ **do**
    if $r_{I_t}(x_t) > 0$ **or** $r_{I_t}(x_t) \leq 0 \cap r_j(x_t) \leq 0$ **then**
      $Q_j(t) \leftarrow Q_j(t-1) + 1$
      $\widehat{\mu}_{j,t} \leftarrow \frac{L_t(\xi_j)}{Q_j(t)} + \left(1 - \frac{1}{Q_j(t)}\right)\widehat{\mu}_{j,t-1}$
    else
      $Q_j(t) \leftarrow Q_j(t-1)$
      $\widehat{\mu}_{j,t} \leftarrow \widehat{\mu}_{j,t-1}$

4.1. UCB-N

Our first result is to extend the UCB-N analysis to show that it admits strong theoretical guarantees even in the setting of time-varying feedback graphs.

Recall that we defined the feedback graph at time t as $G_t = (V, E_t)$, where V is a vertex set of size K representing the experts. We now denote the outgoing edges at each vertex $i \in V_t$ by $N_t^{out}(i)$, where an outgoing edge indicates that playing the source vertex results in feedback for the destination vertex. To distinguish the general setting from the abstention setting, we denote the stochastic losses by l_t, so that $l_t(\xi_{I_t})$ is the loss incurred at time t by playing expert ξ_{I_t}. We let μ_i be the expected loss of expert ξ_t: $\mu_i = \mathbb{E}[l_t(\xi_i)]$, and define $\mu_* = \min_{i\in[K]}\mathbb{E}[l_t(\xi_i)]$ as the loss of the best expert. For convenience, we also use the notation Δ_i to reflect the difference between the expected loss of expert i and the best expert: $\Delta_i = \mu_i - \mu_*$.

The UCB-N algorithm for time-varying feedback graphs is similar to the standard UCB-N of Caron et al. (2012) except that it updates its estimate of the losses according to the varying out-neighborhood $N_t^{out}(I_t)$ at each round t of the chosen expert I_t. Pseudocode of this algorithm is given in Algorithm 7 in the appendix. For arbitrary time-varying feedback graphs, we present the following guarantee:

Theorem 3. *Fix p, and let $\{\mathcal{I}_{t,k}\}_{t\in[T],k\in[p]}$ denote a family*

of subsets of $[K]$ *such that (1) for any fixed* t, $\cup_{k=1}^{p} \mathcal{I}_{t,k} = [K]$ *and (2) for any* $i, j \in \mathcal{I}_{t,k}$, $i \in N_t(j)$.

Then, the regret of UCB-N *after* T *rounds is bounded by:*

$$\sum_{t=1}^{T} \mathbb{E}[l_t(\xi_{I_t}) - l_t(\xi_*)] \leq K\frac{\beta}{\beta-2} + \min_{p,\{\mathcal{I}_{t,k}\}_{t\in[T],k\in[p]}} \sum_{k\in[p]} \left(\max_{\substack{t\in[T]\\ j\in\mathcal{I}_{t,k}}} \Delta_j\right)\left(\max_{\substack{t\in[T]\\ j\in\mathcal{I}_{t,k}}} \frac{8\beta\log(T)}{\Delta_j^2}\right).$$

If our feedback graph structure is fixed throughout all rounds, then we immediately obtain a result that is comparable to Theorem 2 in Caron et al. (2012).

Corollary 1. *If the feedback graph is fixed, i.e.* $G_t = G$, *then the guarantee of Theorem 3 is upper bounded by:*

$$\min_{\mathcal{C}} \sum_{C\in\mathcal{C}} \left(\max_{i\in C} \Delta_i\right)\left(\max_{i\in C} \frac{8\beta\log(T)}{\Delta_i^2}\right) + K\frac{\beta}{\beta-2},$$

where the outer minimum is taken over all clique coverings $\mathcal{C}$ *of the graph* G.

Turning back to the setting of on-line learning with abstention, we again consider the abstention loss $L(\xi)$ and experts ξ of the form $\xi = (h, r)$. We design and analyze the UCB-N.ABS algorithm, which is an application of the general UCB-N algorithm to the abstention scenario. The UCB-N.ABS algorithm calculates a lower confidence bound on the loss of each expert: $\widehat{\mu}_{j,t} - \mathcal{S}_{j,t} = \frac{1}{Q_j(t)}\sum_{s=1}^{t} L_t(\xi_j)1_{I_s=j} - \sqrt{\frac{\beta\log(t)}{2Q_j(t)}}$, where $Q_j(t)$ is the number of times that the algorithm observes the expert $\xi_j \in \mathcal{E}$. Each expert ξ_j is observed when either the chosen expert ξ_{I_t} accepts on a round or when both ξ_j and the chosen expert ξ_{I_t} abstain on a round. Note that the feedback graph determined by x_t at each round can be partitioned into two subsets: the experts that accept and the experts that abstain. For simplicity, define $E_+ = \{i \in [K] : \exists t \in [T], 1_{r_{I_t}(x_t)>0}, 1_{r_i(x_t)>0}\}$ and $E_- = \{i \in [K] : \exists t \in [T], 1_{r_{I_t}(x_t)\leq 0}, 1_{r_i(x_t)\leq 0}\}$ to be the sets of experts that accept and abstain respectively. See Algorithm 1 for the pseudocode of UCB-N.ABS and Appendix B.1 for its accompanying notation. The UCB-N.ABS algorithm admits the following guarantee:

Corollary 2. *After* T *rounds, the regret of* UCB-N.ABS *is at most:*

$$\sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{I_t}) - L_t(\xi_*)] \leq K\frac{\beta}{\beta-2} + \max_{i\in E_+}\Delta_i \max_{i\in E_+}\frac{8\beta\log(T)}{\Delta_i^2} + \max_{i\in E_-}\Delta_i \max_{i\in E_-}\frac{8\beta\log(T)}{\Delta_i^2}.$$

This result highlights how the simple feedback graph structure of the abstention scenario positively affects the regret bound since the resulting sum of the bound is only over two Δ_i. In other words, the dichotomy of the feedback graph based on the sets of experts that accept and abstain is directly reflected in the bound. This is clearly a substantial improvement over standard regret bounds that usually sum over all K experts. Moreover, the regret bounds in Caron et al. (2012) are unable to provide such a guarantee because the feedback graph in our setting changes at every round.

ALGORITHM 2: UCB-D

for $t \geq 1$ **do**
  RECEIVE(x_t)
  $\xi_{I_t} \leftarrow \underset{\xi_j\in\mathcal{E}}{\operatorname{argmin}} \left(\widehat{v}_{j,t-1} - \mathcal{S}_{j,t-1}\right) + c\left(\mathcal{P}_j^-(t-1) - \mathcal{S}_{t-1}\right)$
  if $r_{I_t}(x_t) > 0$ **then**
    Receive(y_t)
  for $\xi_i \in \mathcal{E}$ **do**
    $\mathcal{P}_i^-(t) \leftarrow \frac{1}{t}1_{r_i(x_t)\leq 0} + \left(1 - \frac{1}{t}\right)\mathcal{P}_i^-(t-1)$
    if $r_{I_t}(x_t) > 0$ **or** $r_{I_t}(x_t) \leq 0 \cap r_i(x_t) \leq 0$ **then**
      $Q_i(t) \leftarrow Q_i(t-1) + 1$
      $\widehat{v}_{i,t} \leftarrow \frac{1_{y_t h_i(x_t)\leq 0} 1_{r_i(x)>0}}{Q_i(t)} + \left(1 - \frac{1}{Q_i(t)}\right)\widehat{v}_{i,t-1}$
    else
      $Q_i(t) \leftarrow Q_i(t-1)$
      $\widehat{v}_{i,t} \leftarrow \widehat{v}_{i,t-1}$

4.2. UCB-D

While UCB-N.ABS incorporates additional feedback, its main drawback is that it does not take advantage of the learner's knowledge of the abstention loss. To build upon this algorithm, we design new methods that exploit the structure of the abstention loss. In the following algorithm, which we call UCB-D for *Upper Confidence Bound based on loss Decomposition*, we treat the two contributions to the abstention loss separately. The expected abstention loss can be written as

$$\mathbb{E}[L(\xi, z)] = \bar{c}\,\mathbb{E}[1_{yh(x)\leq 0}1_{r(x)>0}] + c\,\mathbb{P}(r(x) \leq 0),$$

where the last term only depend on the sign of the abstention function r and the first term depends on both h and r. UCB-D separately estimates $\mathbb{P}(r(x) \leq 0)$ and $\mathbb{E}[1_{yh(x)\leq 0}1_{r(x)>0}]$, and then selects the expert that has the smallest empirical expected abstention loss:

$$\xi_{I_t} = \underset{\xi_j\in\mathcal{E}}{\operatorname{argmin}} \left(\widehat{v}_{j,t-1} - \mathcal{S}_{j,t-1}\right) + c\left(\mathcal{P}_j^-(t-1) - \mathcal{S}_{t-1}\right),$$

where $\widehat{v}_{j,t-1}$ approximates $\bar{c}\,\mathbb{E}[1_{yh(x)\leq 0}1_{r(x)>0}]$ and $\mathcal{P}_j^-(t-1)$ estimates the probability that expert j abstains. See Appendix B.1 for the formal definition of the estimates and slacks. Note that because $\mathbb{P}(r(x) \leq 0)$ is an "unsupervised" quantity, we can update our estimate of it at every round, regardless of whether that particular expert accepts

or abstains. Thus, we expect the loss estimate of this algorithm to be superior to that of UCB-N.ABS. The regret of UCB-D is given in the following theorem.

Theorem 4. *On the set of rounds t where the chosen expert abstains, let $\{\mathcal{I}_{t,j}\}_{t\in[T],j\in[2]}$ be a partition of the expert indices $[K]$ into cliques. Then, the regret of* UCB-D *after T rounds is at most:*

$$\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{I_t})-L_t(\xi_*)] \le K\frac{2\beta}{\beta-2} + \max_{i\in E_-}\Delta_i \max_{i\in E_-}\frac{8\beta\log(T)}{\Delta_i^2} + \max_{i\in E_+}\Delta_i \max_{i\in E_+}\frac{8\beta\log(T)}{\Delta_i^2}.$$

While the bound for UCB-D is on the same order as the one for UCB-N.ABS, we expect the actual performance of this algorithm to be better in scenarios where it is often the case that the chosen expert abstains on an example that many of the other experts would not have abstained. In these cases, the estimates in UCB-N.ABS would not be updated. However, UCB-D estimates the probability of abstention separately from the classification error. Thus, the algorithm still learns and refines the loss estimates. We verify this observation in our experiments.

Caron et al. (2012) present tight lower bounds for the case of stochastic bandits with a fixed feedback graph. Because we can again design abstention scenarios with fixed feedback graphs, these bounds also carry over into our setting.

4.3. Sleeping Experts and UCB-H

In this section, we show how the hierarchical structure of the abstention loss lends itself to a "sleeping experts" interpretation of our setting, where not all experts are available at each round.

We first recall the sleeping experts learning scenario. In this framework, at each time t, the environment choses a set $W_t \subseteq \mathcal{E}$ of available (or awake) experts either adversarially or stochastically. Since not all actions are available, the regret is not defined with respect to the best expert in hindsight. Instead, the sleeping regret is characterized by the difference of the loss of the chosen experts and the best available expert according to a fixed ordering determined in hindsight. Specifically, let σ be ordering of the experts $\mathcal{E}$ and define $\sigma(W_t)$ to be the index of the expert with highest rank in σ. The sleeping regret in the stochastic setting is defined as follows:

$$\sum_{t=1}^{T}\mathbb{E}[\ell_t(\xi_{I_t})] - \min_{\sigma}\sum_{t=1}^{T}\mathbb{E}[\ell_t(\xi_{\sigma(W_t)})].$$

A classical algorithm for the sleeping experts problem is the AUER algorithm of Kleinberg et al. (2008), which

ALGORITHM 3: AUER-N

for $t \ge 1$ **do**
  $\xi_{I_t} \leftarrow \operatorname{argmin}_{\xi_j\in W_t}\left\{\widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1}\right\}$
  for $\xi_j \in N^{out}(I_t)$ **do**
    $Q_j(t) \leftarrow Q_j(t-1)+1$
    $\widehat{\mu}_{j,t} \leftarrow \frac{\ell_t(\xi_j)}{Q_j(t)} + \left(1-\frac{1}{Q_j(t)}\right)\widehat{\mu}_{j,t-1}$

chooses the expert with the smallest lowest confidence bound from the available experts and receives its loss. This algorithm is designed only for the standard bandit setting, that is, only the loss of the chosen expert is observed.

Thus, we first extend AUER to take advantage of time-varying feedback graphs, designing an algorithm that we call AUER-N for *AUER with Neighbors*. The idea behind AUER-N is to update the estimate of the expected loss in the out-neighborhood of the chosen expert at every round. Similar to AUER, the expert chosen is the awake pair that has the smallest lowest confidence bound. However, AUER-N also updates the estimate of the expected losses based on the observed feedback. Note that we are assuming for this algorithm and guarantee that the awake set can be chosen adversarially and that $\mu_1 \le \mu_2, \ldots, \le \mu_K$. The pseudocode is given by Algorithm 3, and we present the following sleeping regret bound for AUER-N.

Theorem 5. *For any $T > 0$, the sleeping regret of* AUER-N *is upper bounded by*

$$\sum_{t=1}^{T}\mathbb{E}[\ell_t(\xi_{I_t})] - \min_{\sigma}\sum_{t=1}^{T}\mathbb{E}[\ell_t(\xi_{\sigma(W_t)})] \le O\left(\beta\log T\sum_{j=2}^{K}\frac{1}{\mu_j-\mu_{j-1}}\max_t\frac{T_j(t)}{Q_j(t)}\right).$$

Since $\max_t \frac{T_j(t)}{Q_j(t)} \le 1$ for all experts ξ_j, the sleeping regret bound of AUER-N is always smaller than that of AUER. In particular, if the number of times a learner choses an expert is equal to the number of times that expert was observed, i.e. $T_j(t) = Q_j(t)$, for all arms j, then we recover the sleeping regret bound of AUER.

Since AUER-N does not exploit the hierarchical structure of the abstention loss, we now modify it to the abstention setting and design another algorithm called UCB-H, for *Upper Confidence Bound based on Hierarchical structure*. This algorithm will keep an estimate $\widehat{\nu}_{j,t}$ of the conditional expectation $\mathbb{E}[1_{y h_j(x)\le 0}|r_j(x) > 0]$ along with its corresponding slack term $\widetilde{\mathcal{S}}_{j,t}$. The idea behind UCB-H is that for every point x_t, the algorithm first divides the expert set into accepting experts, for which $r_j(x_t) > 0$, and abstaining experts, for which $r_j(x_t) \le 0$. Then, the algorithm

finds the expert (h_j, r_j) with the lowest confidence bound, that is $\widehat{\nu}_{j,t} - \widetilde{\mathcal{S}}_{j,t}$, over the subset of accepting experts. If this confidence bound is smaller than the abstention cost c, then the algorithm choses (h_j, r_j). If it is higher than c, then it will choose at random an abstaining expert. Algorithm 4 presents the pseudocode, and Appendix B.2 provides a list of additional notation involved in the analysis.

UCB-H can be connected directly to the sleeping expert problem where not all experts are available at each round. Notice that at every time the UCB-H algorithm chooses an expert that abstains, no estimates are updated. Thus, UCB-H can be thought of in terms of sleeping experts by letting the set of awake experts be $W_t = \mathcal{A}_t \cup \{\xi_\diamond\}$, where $\xi_\diamond$ is the all-abstain function with $r_\diamond(x) \leq 0$ for all possible x. In other words at each round, the experts that accept and the all-abstain expert, $\xi_\diamond$, are awake while the rest are sleeping.

Unlike AUER-N, the UCB-H algorithm compares the lowest confidence bound of the available experts with the fixed cost c of the all-abstain function. Moreover, UCB-H estimates the conditional expectation of $\nu_j = \bar{c}\,\mathbb{E}[1_{y h_j(x) \leq 0} | r_j(x) > 0]$ instead of the total expectation. Since the loss of $\xi_\diamond$ is always c, by a slight abuse notation, we let $\nu_{j_\diamond} = c$ and $O_{j_\diamond}(t) = t$ be the number of times the all-abstain expert was observed. By reordering these values such that $\nu_1 \leq \nu_2 \leq \ldots \leq \nu_K$, we can present the following result, a sleeping regret guarantee for UCB-H.

Theorem 6. *For any $T > 0$, the sleeping regret of* UCB-H *is upper bounded by*

$$\sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{I_t})] - \min_{\sigma} \sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{\sigma(W_t)})]$$
$$\leq O\left(\beta \log T \sum_{j=2}^{K} \frac{1}{\nu_j - \nu_{j-1}} \max_t \frac{T_j(t)}{O_j(t)}\right).$$

The lower bounds for sleeping regret in Kleinberg et al. (2008) are based on an adversarial set of awake experts and hence the regret of AUER-N is information theoretically optimal. On the other hand, the regret of UCB-H is not necessarily tight since the awake set is not adversarial.

4.4. Relating sleeping regret to standard regret

It is reasonable to design abstention functions r_j such that $\mathbb{P}(r_j(x) > 0) = \mathbb{P}(r_{j'}(x) > 0)$ for all j, j'. This choice is both natural and also feasible, since the probability of acceptance depends only on the context x and not the label. If we assume that our abstention functions are generated this way, then it follows that the sleeping regret of UCB-H is always at least as good as the sleeping regret bound that one would have obtained without conditional losses.

Theorem 7. *Assume that $\mathbb{P}(r_j(x) > 0) = \mathbb{P}(r_{j'}(x) > 0)$ for all r_j, r'_j, then $\sum_{j=2}^{K} \frac{1}{\nu_j - \nu_{j-1}} \leq \sum_{j=2}^{K} \frac{\mathbb{P}(r_j(x)>0)}{\mu_j - \mu_{j-1}}$. More generally, assume for all $j \in [K]$ such that $\nu_j \neq c$,*

$$\frac{|\,\mathbb{P}(r_j(x) > 0) - \mathbb{P}(r_{j-1}(x) > 0)|}{\mathbb{P}(r_{j-1}(x) > 0)} \leq \frac{\nu_j - \nu_{j-1}}{|\nu_j - c|}.$$

Then, the following inequality holds

$$\sum_{j=2}^{K} \frac{1}{\nu_j - \nu_{j-1}} \leq \sum_{j=2}^{K} \frac{1}{\mu_j - \mu_{j-1}},$$

where $\mu_1 \leq \mu_2 \leq \ldots \leq \mu_K$.

Theorem 7 implies that the sleeping regret bound in terms of the conditional expectation can be much smaller than a sleeping regret bound using the total expectation, such as that of AUER. In particular, if the acceptance probability is such that $\mathbb{P}(r_j > 0) \leq \frac{\mu_j - \mu_{j-1}}{\mu_j - \mu_1}$, then, the the sleeping regret bound of UCB-H is actually comparable to the regret bound of UCB-N.ABS. This scenario is not viable for standard sleeping regret bounds, since it is always the case that $\frac{1}{\mu_j - \mu_1} \leq \frac{1}{\mu_j - \mu_{j-1}}$.

Moreover, if we denote by $\mu_{t,*}$ the expected loss of the best expert that is awake on round t, then as long as $\mu_{t,*} - \mu_* \leq \frac{\log(T)}{T}$, the sleeping regret will be comparable to the standard regret. This includes the case where the best overall expert ξ_* is awake on most of the rounds. In the abstention setting, the experts that are asleep incur a potentially high cost, so we expect this scenario to hold.

Finally, the fact that the abstention option is almost always available in the UCB-H algorithm (unless no expert chooses to abstain on a particular context) means that the algorithm has the ability to *improve* upon the best-in-class expert ξ_*'s abstention function by choosing different accepting experts depending on the awake set. This characteristic is fundamentally different from standard UCB

ALGORITHM 4: UCB-H.

for $t \geq 1$ **do**
  RECEIVE(x_t)
  $\mathbb{A}_t \leftarrow \{\xi_j \in \mathcal{E} : r_j(x_t) > 0\}$
  $\mathbb{R}_t \leftarrow \{\xi_j \in \mathcal{E} : r_j(x_t) \leq 0\}$
  if $\min_{\xi_j \in \mathbb{A}_t} \left\{\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1}\right\} < c$ **then**
    $\xi_{I_t} \leftarrow \operatorname{argmin}_{\xi_j \in \mathbb{A}_t} \left\{\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1}\right\}$
    RECEIVE(y_t)
    for $\xi_j \in \mathbb{A}_t$ **do**
      $O_j(t) \leftarrow O_j(t-1) + 1$
      $\widehat{\nu}_{j,t} \leftarrow \frac{L_t(\xi_j)}{O_j(t)} + \left(1 - \frac{1}{O_j(t)}\right) \widehat{\nu}_{j,t-1}$
  else
    $\xi_{I_t} \leftarrow$UNIFORMSAMPLE($\mathbb{R}_t$)
    $O_j(t) \leftarrow O_j(t-1)$
    $\widehat{\nu}_{j,t} \leftarrow \widehat{\nu}_{j,t-1}$

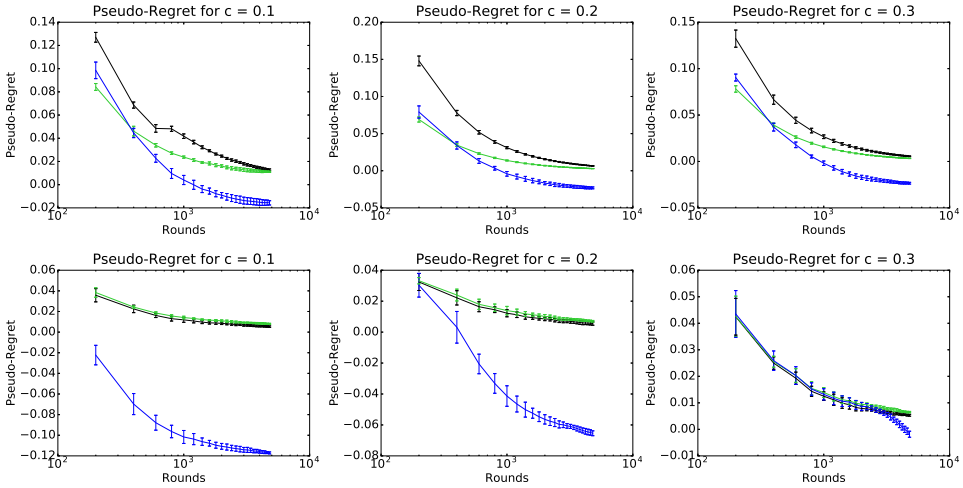

Figure 3: The figure shows the pseudo-regret of UCB-type algorithms with 100 experts as a function of the number of rounds for different values of the cost of abstention. The first row is the results for `Skin` dataset and the second row is the `CIFAR` dataset.UCB-N.ABS is the black line, UCB-D is the green line, and UCB-H is the blue line.

algorithms, which eventually converge on a single expert pair. In the experiments, this benefit will reveal itself in that UCB-H will achieve negative regret on several datasets.

5. Experiments

In this section, we report the results of several experiments comparing the different algorithms analyzed in this paper both on synthetic and real-world datasets. For each dataset, we generated a total of $K = 100$ experts and all the algorithms were tested for a total of $T = 5000$ rounds.

The first real-world dataset is the `Skin` dataset from the UCI data repository. This dataset has three features and the experts used are separating hyperplanes centered at the origin, $\sin(\phi)\cos(\theta)x + \sin(\phi)\sin(\theta)y + \cos(\phi)z = 0$ where $\theta, \phi \in \{\pi/l, 2\pi/l, \ldots, \pi\}$ for $l = 5$. These experts have confidence-based abstention regions around the decision surface with thresholds $\gamma \in \{0.1, 0.2, 0.3, 0.4\}$. The remaining two datasets are the `CIFAR` dataset (Krizhevsky et al., 2009), where we extract the first principal component and use the projection on that as a feature, and a synthetic dataset with two input features. See Appendix C for more details on these datasets and for a description of the types of experts used.

For a fixed set of experts, we first calculated the pseudo-regret by averaging over ten random draws of the data where the best-in-class expert was determined in hindsight as the expert with the minimum average cumulative loss. We then repeated this experiments five times over different sets of experts. For abstention cost $c \in \{0.1, 0.2, 0.3\}$ and misclassification cost $\bar{c} = 1$, Figure 3 above and Figure 5 in Appendix C show the averaged pseudo-regrets with standard deviations for the different UCB-based algorithms as a function of the rounds on the three datasets.

These experimental results show that UCB-H outperforms both UCB-D and UCB-N.ABS in all of the tasks. Note also that UCB-H is the only algorithm that achieves negative regret in these experiments. UCB-D admits a mostly favorable pseudo-regret with respect to UCB-N.ABS. The cost of abstention has an evident effect on all the algorithms especially UCB-H on the `CIFAR` dataset.

6. Conclusion

We introduced and analyzed the setting of on-line learning with an abstention option, showing that efficient learning is possible in both the adversarial and stochastic scenarios. We drew connections between this new on-line setting and existing work involving feedback graphs and sleeping experts. We designed a new algorithm that exploits the structure of the loss function to outperform natural extensions of existing algorithms. Remarkably, this algorithm even achieves negative regret in our experiments. In the process, we also presented a new analysis for stochastic multi-armed bandits with time-varying feedback graphs along with a new algorithm and analysis for sleeping stochastic multi-armed bandits with time-varying feedback graphs, generalizing and extending prior work in both cases. The concept of on-line learning with abstention is general. This work can be extended to similar problems, including the multi-class and regression settings, as well as more realistic scenarios, such as on-line learning with a budget constraint.

References

Alon, Noga, Cesa-Bianchi, Nicolò, Gentile, Claudio, and Mansour, Yishay. From bandits to experts: A tale of dom- ination and independence. In *NIPS*, 2013.

Alon, Noga, Cesa-Bianchi, Nicolò, Gentile, Claudio, Mannor, Shie, Mansour, Yishay, and Shamir, Ohad. Nonstochastic multi-armed bandits with graph-structured feedback. In *CoRR*, 2014.

Alon, Noga, Cesa-Bianchi, Nicolò, Dekel, Ofer, and Koren, Tomer. Online learning with feedback graphs: Beyond bandits. In *JMLR*, pp. 23–35, 2015.

Bartlett, Peter and Wegkamp, Marten. Classification with a reject option using a hinge loss. *JMLR*, pp. 291–307, 2008.

Caron, Stephane, Kveton, Branislav, Lelarge, Marc, and Bhagat, Smriti. Leveraging side observations in stochastic bandits. In *UAI*, 2012.

Cesa-Bianchi, Nicolo, Mansour, Yishay, and Stoltz, Gilles. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2-3):321–352, 2007.

Cohen, Alon, Hazan, Tamir, and Koren, Tomer. Online learning with feedback graphs without the graphs. *arXiv preprint arXiv:1605.07018*, 2016.

Cortes, Corinna, DeSalvo, Giulia, and Mohri, Mehryar. Learning with rejection. In *ALT*, pp. 67–82. Springer, Heidelberg, Germany, 2016a.

Cortes, Corinna, DeSalvo, Giulia, and Mohri, Mehryar. Boosting with abstention. In *NIPS*. MIT Press, 2016b.

Kleinberg, Robert, Niculescu-Mizil, Alexandru, and Sharma, Yogeshwer. Regret bounds for sleeping experts and bandits. *Machine Learning*, 2008.

Kocák, Tomáš, Neu, Gergely, Valko, Michal, and Munos, Rémi. Efficient learning by implicit exploration in bandit problems with side observations. In *NIPS*, pp. 613–621, 2014.

Krizhevsky, Alex, Nair, Vinod, and Hinton, Geoffrey. Cifar-10 (canadian institute for advanced research), 2009. URL http://www.cs.toronto.edu/~kriz/cifar.html.

Li, Lihong, Littman, Michael, and Thomas, Walsh. Knows what it knows: A framework for self-aware learning. In *ICML*, 2008.

Littlestone, Nick and Warmuth, Manfred K. The weighted majority algorithm. *Information and computation*, 108 (2):212–261, 1994.

Mannor, Shie and Shamir, Ohad. From bandits to experts: On the value of side-observations. *NIPS*, pp. 291–307, 2011.

Neu, Gergely. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. In *NIPS*, pp. 3168–3176, 2015.

Sayedi, Amin, Zadimoghaddam, Morteza, and Blum, Avrim. Trading off mistakes and dont-know predictions. In *NIPS*, 2010.

Zhang, Chicheng and Chaudhuri, Kamalika. The extended littlestone's dimension for learning with mistakes and abstentions. In *COLT*, 2016.

ALGORITHM 5: EXP3.ABS adapted from Alon et al. (2015)

Require: exploration rate $\gamma \in [0,1]$, learning rate $\eta > 0$, u uniform distribution, q_1 uniform distribution
for $t \leftarrow 1, 2, \ldots$ **do**
  RECEIVE(x_t)
  $p_t \leftarrow (1-\gamma)q_t + \gamma u$
  $\xi_{I_t} \leftarrow$ SAMPLE(p_t)
  if $r_{I_t}(x_t) > 0$ **then**
    RECEIVE(y_t)
  RECEIVE($L_t(\xi_{I_t})$)
  $P_t(\xi) \leftarrow 1_{r(x_t)>0} \sum_{\tilde{\xi}\in\mathcal{E}:\tilde{r}(x_t)>0} p_t(\tilde{\xi}) + 1_{r(x_t)\leq 0}$
  $\widehat{l}_t(\xi) \leftarrow \frac{L_t(\xi)}{P_t(\xi)} \left(1_{r_{I_t}(x_t)>0} 1_{r(x_t)>0} + 1_{r_{I_t}(x_t)\leq 0}\right)$
  $q_{t+1}(\xi) \leftarrow \frac{q_t(\xi)\exp(-\eta \widehat{l}_t(\xi))}{\sum_{\tilde{\xi}\in\mathcal{E}} q_t(\tilde{\xi})\exp(-\eta \widehat{l}_t(\tilde{\xi}))}$

ALGORITHM 6: EXP3-IX.ABS adapted from Alon et al. (2015)

Require: learning rate $\eta > 0$, p_1 uniform distribution, implicit exploration rate $\kappa \in [0,1]$
for $t \leftarrow 1, 2, \ldots$ **do**
  RECEIVE(x_t)
  $\xi_{I_t} \leftarrow$ SAMPLE(p_t)
  if $r_t(x_t) > 0$ **then**
    RECEIVE(y_t)
  RECEIVE($L_t(\xi_{I_t})$).
  $P_t(\xi) \leftarrow 1_{r(x_t)>0} \sum_{\tilde{\xi}\in\mathcal{E}:\tilde{r}(x_t)>0} p_t(\tilde{\xi}) + 1_{r(x_t)\leq 0}$
  $\tilde{L}_t(\xi) \leftarrow \frac{L_t(\xi)}{P_t(\xi)+\kappa} \left(1_{r_{I_t}(x_t)>0} 1_{r(x_t)>0} + 1_{r_{I_t}(x_t)\leq 0}\right)$
  $p_{t+1}(\xi) \leftarrow \frac{p_t(\xi)\exp(-\eta \tilde{L}_t(\xi))}{\sum_{\tilde{\xi}\in\mathcal{E}} p_t(\tilde{\xi})\exp(-\eta \tilde{L}_t(\tilde{\xi}))}$

A. Additional material for the adversarial setting

A.1. Pseudocode

Algorithm 5 presents the pseudocode for EXP3.ABS, an algorithm for on-line learning with abstention under an adversarial data model that guarantees small pseudo-regret.

Algorithm 6 presents the pseudocode for EXP3-IX.ABS, an algorithm for on-line learning with abstention under an adversarial data model that guarantees small regret with high probability.

The theoretical guarantees for both algorithms are presented and derived in Appendix A.2.

A.2. Proof of theorems

This section contains the proof of the theorems of the adversarial setting of Section 3. We first prove the guarantee of EXP3.ABS and then that of EXP3-IX.ABS.

Theorem 1. *For any $T \geq 1$ and any $\xi^\star \in \mathcal{E}$, EXP3.ABS admits the guarantee:*

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)L_t(\xi)\right] - \sum_{t=1}^{T} L_t(\xi^\star) \leq \gamma T + \frac{\log K}{\eta} + \eta\sum_{t=1}^{T}\left(8c^2 \log\frac{4KK_{abs,t}}{\gamma} + 8\bar{c}^2 \log\frac{4KK_{acc,t}}{\gamma}\right).$$

Proof. By applying the standard second-order regret bound of Hedge in (Cesa-Bianchi et al., 2007) to distributions $q_1, \ldots, q_t$ generated by EXP3.ABS and to the non-negative loss estimates $\widehat{l}_t(\xi)$, the following holds

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} q_t(\xi)\,\mathbb{E}\left[\widehat{l}_t(\xi)\right] - \sum_{t=1}^{T}\mathbb{E}\left[\widehat{l}_t(\xi^\star)\right]\right] \leq \frac{\log K}{\eta} + \eta\sum_{t=1}^{T}\mathbb{E}\left[\sum_{\xi\in\mathcal{E}} q_t(\xi)\,\mathbb{E}\left[\widehat{l}_t(\xi)^2\right]\right]$$

for any fixed $(\xi^\star) \in \mathcal{E}$. Using the fact that $\mathbb{E}\left[\widehat{l}_t(\xi)\right] = L_t(\xi)$ and $\mathbb{E}\left[\widehat{l}_t(\xi)^2\right] = \frac{L_t(\xi)^2}{P_t(\xi)}$,

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} q_t(\xi)L_t(\xi) - \sum_{t=1}^{T} L_t(\xi^\star)\right] \leq \frac{\log K}{\eta} + \eta\sum_{t=1}^{T}\mathbb{E}\left[\sum_{\xi\in\mathcal{E}}\frac{q_t(\xi)}{P_t(\xi)}L_t(\xi)^2\right].$$

Now for each t, we can split the graph in V into two graphs $V_{abs,t}$ and $V_{acc,t}$ where if a node ξ_i is abstaining then $\xi_i \in V_{abs,t}$ and if a node ξ_j is accepting, then $\xi_j \in V_{acc,t}$. Thus, we can write the following

$$\begin{aligned}
&\sum_{\xi\in\mathcal{E}}\frac{q_t(\xi)}{P_t(\xi)}L_t(\xi)^2 \\
&= \sum_{\xi\in V_{abs,t}}\frac{q_t(\xi)}{P_t(\xi)}L_t(\xi)^2 \\
&+ \sum_{\xi\in V_{acc,t}}\frac{q_t(\xi)}{P_t(\xi)}L_t(\xi)^2 \\
&\leq \sum_{\xi\in V_{abs,t}}\frac{q_t(\xi)}{P_t(\xi)}c^2 + \sum_{\xi\in V_{acc,t}}\frac{q_t(\xi)}{P_t(\xi)}\bar{c}^2.
\end{aligned}$$

The inequality holds since for the abstained points at time t, we know that the $L(\xi) = c$ while for the accepted points, we know that $L(\xi) \leq \bar{c}$. Since all nodes $\xi \in V_{abs,t}$ have self-loops and $p_t(\xi) \geq \frac{\gamma}{K}$ because we mixed with the uniform distribution, we can apply Lemma 5 in (Alon et al.,

2015) with $\epsilon = \frac{\gamma}{K}$. Thus, we have the following inequalities

$$\sum_{\xi \in V_{abs,t}} \frac{q_t(\xi)}{P_t(\xi)} c^2 \leq 2 \sum_{\xi \in V_{abs,t}} \frac{p_t(\xi)}{P_t(\xi)} c^2 \leq 8c^2 \alpha_{abs,t} \log \frac{4KK_{abs,t}}{\gamma \alpha_{abs,t}}$$

where the first inequality is due to the fact that $p_t(\xi) \geq (1-\gamma) q_t(\xi) \geq \frac{1}{2} q_t(\xi)$ and where $\alpha_{abs,t} = \alpha(V_{abs,t})$ is the independence number of $V_{abs,t}$. Similarly,

$$\sum_{\xi \in V_{acc,t}} \frac{q_t(\xi)}{P_t(\xi)} \bar{c}^2 \leq 2 \sum_{\xi \in V_{acc,t}} \frac{p_t(\xi)}{P_t(\xi)} \bar{c}^2 \leq 8 \alpha_{acc,t} \bar{c}^2 \log \frac{4KK_{acc,t}}{\gamma \alpha_{acc,t}}$$

where $\alpha_{acc,t} = \alpha(V_{acc,t})$ is the independence number of $V_{acc,t}$. Now, for the feedback graphs in this setting, the $\alpha_{abs,t} = \alpha_{acc,t} = 1$ for all t, which is the best we can hope for. Thus,

$$\sum_{\xi \in V_{abs,t}} \frac{q_t(\xi)}{P_t(\xi)} c^2 + \sum_{\xi \in V_{acc,t}} \frac{q_t(\xi)}{P_t(\xi)} \bar{c}^2 \leq 8c^2 \log \frac{4KK_{abs,t}}{\gamma} + 8\bar{c}^2 \log \frac{4KK_{acc,t}}{\gamma}.$$

Using then the fact that

$$\sum_{\xi \in V} p_t(\xi) L_t(\xi) \leq \sum_{\xi \in V} q_t(\xi) L_t(\xi) + \gamma,$$

the regret bound in this case can be written as

$$\mathbb{E}\left[\sum_{t=1}^T \sum_{\xi \in V} p_t(\xi) L_t(\xi)\right] - \sum_{t=1}^T L_t(\xi^\star) \leq \gamma T + \frac{\log K}{\eta} + \eta \sum_{t=1}^T \left[8c^2 \log \frac{4KK_{abs,t}}{\gamma} + 8\bar{c}^2 \log \frac{4KK_{acc,t}}{\gamma}\right].$$

□

Theorem 2. *For any $\delta > 0$, $\gamma \leq \frac{1}{\bar{c}K}$ and $\gamma' \leq \frac{1}{cK}$,* EXP3-IX.ABS *admits the following guarantee with probability at least $1 - 4\delta$:*

$$\sum_{t=1}^T L_t(\xi_{I_t}) - L_t(\xi^\star) \leq \max\{c, \bar{c}\} \sqrt{\frac{T \log(1/\delta)}{2}} + \left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right) \sum_{t=1}^T 2\bar{c} f(K_{acc,t}) + \left(\frac{\eta}{2} c + \kappa + (e-2)\gamma'\right) \sum_{t=1}^T 2c f(K_{abs,t}) + \frac{\log K}{\eta} + \frac{3\eta}{4} \frac{K \log\left(\frac{K}{\delta}\right)}{\kappa} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'},$$

where $f(x) = \log\left(1 + \left\lceil \frac{x^2}{\kappa} \right\rceil + x\right) + 1$.

Proof. We want to relate the losses that we incur, $L_t(\xi_t)$ to the surrogate losses $\tilde{L}_t(\xi_t)$. We do this in two steps. First, note that by applying Hoeffding's inequality,

$$\sum_{t=1}^T L_t(\xi_t) \leq \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) L_t(\xi) + \max\{c, \bar{c}\} \sqrt{\frac{T \log(1/\delta)}{2}}.$$

Next, we can relate the expected loss to the expected surrogate loss:

$$\begin{aligned}
&\sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) L_t(\xi) \\
&= \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) \tilde{L}_t(\xi) - p_t(\xi) \tilde{L}_t(\xi) + p_t(\xi) L_t(\xi) \\
&= \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) \tilde{L}_t(\xi) \\
&\quad - \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) L_t(e\xi) \left[\frac{1_{\xi \in O_t}}{P_t(\xi) + \kappa} - \frac{P_t(\xi) + \kappa}{P_t(\xi) + \kappa}\right] \\
&= \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) \tilde{L}_t(\xi) \\
&\quad + \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} p_t(\xi) L_t(\xi) \left[\frac{P_t(\xi) - 1_{\xi \in O_t}}{P_t(\xi) + \kappa}\right] \\
&\quad + \sum_{t=1}^T \sum_{\xi \in \mathcal{E}} \frac{\kappa p_t(\xi) L_t(\xi)}{P_t(\xi) + \kappa}.
\end{aligned}$$

Now, since the algorithm plays a weighted majority algorithm on $\tilde{L}_t$, we can bound the cumulative expected surro-

gate loss by:

$$\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)\tilde{L}_t(\xi) \leq \sum_{t=1}^{T}\tilde{L}_t(\xi^\star) + \frac{\log K}{\eta} + \frac{\eta}{2}\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)\tilde{L}_t(\xi)^2.$$

By definition of $\tilde{L}_t$, we can upper bound the last term as follows:

$$\begin{aligned}
&\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)\tilde{L}_t(\xi)^2 \\
&= \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)\left[1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}} + 1_{r(x_t)<0c}\right]}{P_t(\xi)+\kappa}\tilde{L}_t(\xi) \\
&= \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi) \\
&\quad + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r(x_t)<0c}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi).
\end{aligned}$$

We can bound each component in this last term in the following way:

$$\begin{aligned}
&\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi) \\
&= \frac{\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}1_{\xi\in O_t}L_t(\xi)}{P_t(\xi)+\kappa} \\
&\leq \frac{\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}1_{\xi\in O_t}L_t(\xi)}{P_t(\xi)+\kappa\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}L_t(\xi)} \\
&= \frac{\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}1_{\xi\in O_t}\frac{L_t(\xi)}{P_t(\xi)}}{1+\kappa\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\frac{L_t(\xi)}{P_t(\xi)}} \\
&\leq \frac{1}{2\kappa}\log\left(1+2\kappa\frac{L_t(\xi)1_{\xi\in O_t}}{P_t(\xi)}\right).
\end{aligned}$$

The last line follows from the fact that $\frac{x}{1+(x/2)} \leq \log(1+x)$ for $x \geq 0$.

Thus, it follows that

$$\begin{aligned}
&\exp\left(2\kappa\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi)\right) \\
&\leq 1 + 2\kappa\frac{L_t(\xi)1_{\xi\in O_t}}{P_t(\xi)}.
\end{aligned}$$

Let $(\mathcal{F}_t)_{t=1}^{\infty}$ be a filtration such that $\mathcal{F}_t$ is the history of the game up to knowledge of x_t. Then we can write

$$\begin{aligned}
&\mathbb{E}\left[\exp\left(2\kappa\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi)\right)|\mathcal{F}_t\right] \\
&\leq \mathbb{E}\left[1 + 2\kappa\frac{L_t(\xi)1_{\xi\in O_t}}{P_t(\xi)}|\mathcal{F}_t\right] \\
&\leq 1 + 2\kappa L_t(\xi) \\
&\leq \exp\left(2\kappa L_t(\xi)\right).
\end{aligned}$$

The last line follows from the fact that $1 + x \leq e^x$ for all $x \in \mathbb{R}$.

Thus, we can use this exponential bound to write that

$$\begin{aligned}
&\mathbb{P}\left[\sum_{t=1}^{T} 2\kappa\frac{p_t(\xi)}{P_t(\xi)+\kappa}1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}\right. \\
&\quad \left.\left(\tilde{L}_t(\xi) - L_t(\xi)\right) > \frac{\log\left(\frac{K}{\delta}\right)}{2\kappa}\right] \\
&\leq \mathbb{E}\left[\exp\left(2\kappa\frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi)\right)\right]\frac{\delta}{K} \\
&\leq \frac{\delta}{K}.
\end{aligned}$$

Similarly, we can show that

$$\mathbb{P}\left[\sum_{t=1}^{T} 2\kappa\frac{p_t(\xi)}{P_t(\xi)+\kappa}1_{r_t(x_t)>0}1_{r(x_t)<0c} \left(\tilde{L}_t(\xi) - L_t(\xi)\right) > \frac{\log\left(\frac{K}{\delta}\right)}{2\kappa}\right] \leq \frac{\delta}{K}.$$

By taking a union bound over all $\xi \in \mathcal{E}$, it now follows that with probability at least $1 - 2\delta$,

$$\begin{aligned}
&\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi) \\
&\quad + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r(x_t)<0c}}{P_t(\xi)+\kappa}\tilde{L}_t(\xi) \\
&= \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}}}{P_t(\xi)+\kappa}L_t(\xi) + \\
&\quad + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r(x_t)<0c}}{P_t(\xi)+\kappa}L_t(\xi) + \frac{K\log\left(\frac{K}{\delta}\right)}{\kappa} \\
&\leq \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0\bar{c}^2}}{P_t(\xi)+\kappa} + \\
&\quad + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r(x_t)<0c^2}}{P_t(\xi)+\kappa} + \frac{K\log\left(\frac{K}{\delta}\right)}{\kappa}.
\end{aligned}$$

Putting all the pieces together, we can now say that

$$\begin{aligned}\sum_{t=1}^{T} L_t(\xi_t) \leq \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} \\ + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)L_t(\xi)\left[\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right] \\ + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{\kappa p_t(\xi)L_t(\xi)}{P_t(\xi)+\kappa} + \sum_{t=1}^{T}\tilde{L}_t(e^*) + \frac{\log K}{\eta} \\ + \frac{\eta}{2}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r_t(x_t)>0}1_{r(x_t)>0}\bar{c}^2}{P_t(\xi)+\kappa} + \right. \\ \left. + \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi)1_{r(x_t)<0}c^2}{P_t(\xi)+\kappa} + \frac{K\log\left(\frac{K}{\delta}\right)}{\kappa}\right].\end{aligned}$$

Now consider the second term in the upper bound: $\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} p_t(\xi)L_t(\xi)\left[\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right]$. We can decompose this as follows:

$$\begin{aligned}&\sum_{\xi\in\mathcal{E}} p_t(\xi)L_t(\xi)\left[\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right] \\ &= \sum_{\xi\in\mathcal{E}} p_t(\xi)\bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\left[\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right] \\ &\quad + \sum_{\xi\in\mathcal{E}} p_t(\xi)c1_{r(x_t)<0}\left[\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right] \\ &= X_t(\xi) + Y_t(\xi).\end{aligned}$$

Notice that each of $X_t(\xi)$ and $Y_t(\xi)$ is a martingale-difference with $|X_t| \leq K\bar{c}$ and $|Y_t| \leq Kc$.

Moreover, we can bound the conditional variance as:

$$\begin{aligned}&\mathbb{E}\left[X_t(\xi)^2|\mathcal{F}_t\right] \\ &= \mathbb{E}\left[\left(\sum_{\xi\in\mathcal{E}} p_t(\xi)\bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{P_t(\xi)-1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right)^2\middle|\mathcal{F}_t\right] \\ &\leq \mathbb{E}\left[\left(\sum_{\xi\in\mathcal{E}} p_t(\xi)\bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\left[\frac{1_{\xi\in O_t}}{P_t(\xi)+\kappa}\right]\right)^2\middle|\mathcal{F}_t\right] \\ &= \mathbb{E}\left[\sum_{\xi,\xi'\in\mathcal{E}} \bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{p_t(\xi)p_t(\xi')1_{\xi\in O_t}1_{\xi'\in O_t}}{(P_t(\xi)+\kappa)(P_t(\xi')+\kappa)}\middle|\mathcal{F}_t\right] \\ &\leq \mathbb{E}\left[\sum_{\xi,\xi'\in\mathcal{E}} \bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{p_t(\xi)p_t(\xi')1_{\xi\in O_t}}{(P_t(\xi)+\kappa)(P_t(\xi')+\kappa)}\middle|\mathcal{F}_t\right] \\ &= \sum_{\xi,\xi'\in\mathcal{E}} \bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{p_t(\xi)p_t(\xi')P_t(\xi)}{(P_t(\xi)+\kappa)(P_t(\xi')+\kappa)} \\ &\leq \sum_{\xi,\xi'\in\mathcal{E}} p_t(\xi)\bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{p_t(\xi')}{P_t(\xi')+\kappa}.\end{aligned}$$

By Freedman's inequality, it follows that with probability at least $1-\delta$, for any $\gamma \leq \frac{1}{K\bar{c}}$,

$$\begin{aligned}\sum_{t=1}^{T} X_t(\xi) \leq \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} \\ + (e-2)\gamma\sum_{t=1}^{T}\sum_{\xi'\in\mathcal{E}} \bar{c}1_{r_t(x_t)>0}1_{r(x_t)>0}\frac{p_t(\xi')}{P_t(\xi')+\kappa}.\end{aligned}$$

We can bound term Y_t similarly to say that with probability at least $1-\delta$, for any $\gamma' \leq \frac{1}{Kc}$,

$$\begin{aligned}\sum_{t=1}^{T} Y_t(\xi) \leq \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'} \\ + (e-2)\gamma'\sum_{t=1}^{T}\sum_{e'\in\mathcal{E}} c1_{r(x_t)<0}\frac{p_t(\xi')}{P_t(\xi')+\kappa}.\end{aligned}$$

It now follows that with probability at least $1-4\delta$, for any

$\gamma \leq \frac{1}{K\bar{c}}$ and $\gamma' \leq \frac{1}{Kc}$,

$$\begin{aligned}
\sum_{t=1}^{T} L_t(\xi_t) &\leq \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} \\
&+ \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{\kappa p_t(\xi) L_t(\xi)}{P_t(\xi)+\kappa} + \sum_{t=1}^{T} \tilde{L}_t(\xi^\star) + \frac{\log K}{\eta} \\
&+ \frac{\eta}{2}\left[\sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r_t(x_t)>0} 1_{r(x_t)>0}\bar{c}^2}{P_t(\xi)+\kappa}\right. \\
&\left.+ \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r(x_t)<0} c^2}{P_t(\xi)+\kappa} + \frac{K\log\left(\frac{K}{\delta}\right)}{\kappa}\right] \\
&+ \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'} \\
&+ (e-2)\gamma \sum_{t=1}^{T}\sum_{\xi'\in\mathcal{E}} \bar{c} 1_{r_t(x_t)>0} 1_{r(x_t)>0} \frac{p_t(\xi')}{P_t(\xi')+\kappa} \\
&+ (e-2)\gamma' \sum_{t=1}^{T}\sum_{\xi'\in\mathcal{E}} c 1_{r(x_t)<0} \frac{p_t(\xi')}{P_t(\xi')+\kappa} \\
&= \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} + \sum_{t=1}^{T} \tilde{L}_t(e^*) + \frac{\log K}{\eta} \\
&+ \left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right) \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r_t(x_t)>0} 1_{r(x_t)>0}\bar{c}}{P_t(\xi)+\kappa} \\
&+ \left(\frac{\eta}{2}c + \kappa + (e-2)\gamma'\right) \sum_{t=1}^{T}\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r(x_t)<0} c}{P_t(\xi)+\kappa} \\
&+ \frac{\eta}{2}\frac{K\log\left(\frac{K}{\delta}\right)}{\kappa} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'}.
\end{aligned}$$

By Lemma 1 in (Kocák et al., 2014),

$$\begin{aligned}
\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r_t(x_t)>0} 1_{r(x_t)>0}\bar{c}}{P_t(\xi)+\kappa} &\leq \sum_{\xi\in\mathcal{E}_{\text{acc},t}} \frac{p_t(\xi)\bar{c}}{P_t(\xi)+\kappa} \\
&\leq 2\bar{c}\left[\log\left(1 + \left\lceil\frac{K_{\text{acc},t}^2}{\kappa}\right\rceil + K_{\text{acc},t}\right) + 1\right],
\end{aligned}$$

and

$$\begin{aligned}
\sum_{\xi\in\mathcal{E}} \frac{p_t(\xi) 1_{r(x_t)<0} c}{P_t(\xi)+\kappa} &\leq \sum_{\xi\in\mathcal{E}_{\text{abs},t}} \frac{p_t(\xi) c}{P_t(\xi)+\kappa} \\
&\leq 2c\left[\log\left(1 + \left\lceil\frac{K_{\text{abs},t}^2}{\kappa}\right\rceil + K_{\text{abs},t}\right) + 1\right].
\end{aligned}$$

Thus, we have shown that

$$\begin{aligned}
\sum_{t=1}^{T} L_t(\xi_t) &\leq \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} \\
&+ \sum_{t=1}^{T} \tilde{L}_t(\xi^\star) + \frac{\log K}{\eta} \\
&+ \left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right) \\
&\sum_{t=1}^{T} 2\bar{c}\left[\log\left(1 + \left\lceil\frac{K_{\text{acc},t}^2}{\kappa}\right\rceil + K_{\text{acc},t}\right) + 1\right] \\
&+ \left(\frac{\eta}{2}c + \kappa + (e-2)\gamma'\right) \\
&\sum_{t=1}^{T} 2c\left[\log\left(1 + \left\lceil\frac{K_{\text{abs},t}^2}{\kappa}\right\rceil + K_{\text{abs},t}\right) + 1\right] \\
&+ \frac{\eta}{2}\frac{K\log\left(\frac{K}{\delta}\right)}{\kappa} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'} \\
&\leq \max\{c,\bar{c}\}\sqrt{\frac{T\log(1/\delta)}{2}} + \sum_{t=1}^{T} L_t(e^*) + \frac{\log K}{\eta} \\
&+ \left(\frac{\eta}{2}\bar{c} + \kappa + (e-2)\gamma\right) \\
&\sum_{t=1}^{T} 2\bar{c}\left[\log\left(1 + \left\lceil\frac{K_{\text{acc},t}^2}{\kappa}\right\rceil + K_{\text{acc},t}\right) + 1\right] \\
&+ \left(\frac{\eta}{2}c + \kappa + (e-2)\gamma'\right) \\
&\sum_{t=1}^{T} 2c\left[\log\left(1 + \left\lceil\frac{K_{\text{abs},t}^2}{\kappa}\right\rceil + K_{\text{abs},t}\right) + 1\right] \\
&+ \frac{3\eta}{4}\frac{K\log\left(\frac{K}{\delta}\right)}{\kappa} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma} + \frac{\log\left(\frac{1}{\delta}\right)}{\gamma'}.
\end{aligned}$$

□

B. Additional material for the stochastic setting

B.1. Definitions

In this section, we provide definitions for UCB-N and AUER-N as well as the pseudocode of UCB-N algorithm by Caron et al. (2012) for completeness in Algorithm 7. We also include tables of the definitions needed for UCB-N.ABS, UCB-D, and UCB-H. Note that for example $Q_j(t)$ in the definition of UCB-N.ABS and UCB-D is an instantiation to the abstention setting of the $Q_j(t)$ defined for UCB-N and AUER-N. In other words, for the abstention setting, the expert j is in out-neighborhood of the chosen expert whenever $r_{I_s}(x_s) \leq 0 \cap r_j(x_s) \leq 0$ or $r_{I_s}(x_s) > 0$. Similarly for $\hat{\mu}_{j,t}$ and $\mathcal{S}_{j,t}$ are instantiations to the abstention setting.

ALGORITHM 7: UCB-N

for $t \geq 1$ **do**
  $\xi_{I_t} \leftarrow \operatorname{argmin}_{\xi_j \in \mathcal{E}} \left\{ \widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \right\}$
  for $\xi_j \in N^{out}(I_t)$ **do**
    $Q_j(t) \leftarrow Q_j(t-1) + 1$
    $\widehat{\mu}_{j,t} \leftarrow \frac{\ell_t(\xi_j)}{Q_j(t)} + \left(1 - \frac{1}{Q_j(t)}\right) \widehat{\mu}_{j,t-1}$

Definitions for UCB-N and AUER-N.

$$T_j(t) = \sum_{s=1}^{t} 1_{I_s = j}$$

$$Q_j(t) = \sum_{s=1}^{t} 1_{j \in N^{out}(I_s)}$$

$$\widehat{\mu}_{j,t} = \frac{1}{Q_j(t)} \sum_{s=1}^{t} \ell_s(\xi_j) 1_{j \in N^{out}(I_s)}$$

$$\mathcal{S}_{j,t} = \sqrt{\frac{\beta \log t}{2 Q_j(t)}}$$

Definitions for UCB-N.ABS and UCB-D.

$$Q_j(t) = \sum_{s=1}^{t} \left(1_{r_{I_s}(x_s) \leq 0} 1_{r_j(x_s) \leq 0} + 1_{r_{I_s}(x_s) > 0}\right)$$

$$\widehat{\mu}_{j,t} = \frac{1}{Q_j(t)} \sum_{s=1}^{t} L_s(\xi_j) \big(1_{r_{I_s}(x_s) \leq 0} 1_{r(x_s) \leq 0} + 1_{r_{I_s}(x_s) > 0}\big)$$

$$\mathcal{S}_{j,t} = \sqrt{\frac{\beta \log(t)}{2 Q_j(t)}}$$

$$\mathcal{P}_j^-(t) = \frac{1}{t} \sum_{s=1}^{t} 1_{r_j(x_s) \leq 0}$$

$$\widehat{v}_{j,t} = \frac{1}{Q_j(t)} \sum_{s=1}^{t} 1_{y_s h_j(x_s) \leq 0} 1_{r_j(x_s) > 0} \big(1_{r_{I_s}(x_s) \leq 0} 1_{r(x_s) \leq 0} + 1_{r_{I_s}(x_s) > 0}\big)$$

$$\mathcal{S}_t = \sqrt{\frac{\beta \log t}{2t}}$$

Definitions for UCB-H.

$$O_j(t) = \sum_{s=1}^{t} 1_{r_j(x_s) > 0} 1_{r_{I_s}(x_s) > 0}$$

$$\widehat{\nu}_{j,t} = \frac{1}{O_j(t)} \sum_{s=1}^{t} 1_{y_s h_j(x_s) \leq 0} 1_{r_j(x_s) > 0} 1_{r_{I_s}(x_s) > 0}$$

$$\widetilde{\mathcal{S}}_{j,t} = \sqrt{\frac{\beta \log t}{2 O_j(t)}}$$

B.2. Proofs of theorems

In this section, we prove the theorems in the stochastic setting of Section 4 by first focusing on the standard regret guarantees and then turning to the sleeping regret bounds.

B.2.1. STANDARD REGRET

We now prove the general theorem of UCB-N based on cliques structure of the feedback graphs.

Theorem 3. *Fix p, let $\{\mathcal{I}_{t,k}\}_{t \in [T], k \in [p]}$ denote a family of subsets of $[K]$ such that (1) for any fixed t, $\cup_{k=1}^{p} \mathcal{I}_{t,k} = [K]$ and (2) for any $i, j \in \mathcal{I}_{t,k}$, $i \in N_t(j)$.*

Then, the regret of UCB-N *after T rounds is bounded by:*

$$\sum_{t=1}^{T} \mathbb{E}[l_t(\xi_{I_t}) - l_t(\xi_*)] \leq \min_{p} \min_{\{\mathcal{I}_{t,k}\}_{t \in [T], k \in [p]}} \sum_{k \in [p]} \left(\max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} \Delta_j \right) \left(\max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} \frac{8\beta \log(T)}{\Delta_j^2} \right) + K \frac{\beta}{\beta - 2}.$$

Proof.

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[l_t(\xi_{I_t}) - l_t(\xi_*)] \\
&= \sum_{i=1}^{K} \sum_{t=1}^{T} \mathbb{E}[l_t(\xi_i) - l_t(\xi_*)] \, \mathbb{E}[1_{I_t = i}] \\
&= \sum_{i=1}^{K} \sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t = i}] \\
&= \sum_{i=1}^{K} \sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t = i} (1_{Q_i(t-1) \leq s_i} + 1_{Q_i(t-1) > s_i})],
\end{aligned}$$

where $Q_i(t) = \sum_{s=1}^{t} 1_{i \in N_t^{out}(I_t)}$ and s_i is to be determined.

We can decompose this sum and bound the latter term as follows:

$$\begin{aligned}
&\mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)>s_i}] \\
&= \mathbb{P}[I_t = i, Q_i(t-1) > s_i] \\
&\leq \mathbb{P}[\widehat{\mu}_{i,t-1} - \mathcal{S}_{i,t} \leq \widehat{\mu}_{*,t} - \mathcal{S}_{*,t}, Q_i(t-1) > s_i].
\end{aligned}$$

We can rewrite the first condition as:

$$\begin{aligned}
0 &\leq \widehat{\mu}_{*,t-1} - \mathcal{S}_{*,t-1} - \widehat{\mu}_{i,t-1} + \mathcal{S}_{i,t-1} \\
0 &\leq (\widehat{\mu}_{*,t-1} - \mathcal{S}_{*,t-1} - \mu_*) \\
&\quad + (\mu_i - \widehat{\mu}_{i,t-1} + \mathcal{S}_{i,t-1} - 2\mathcal{S}_{i,t-1}) \\
&\quad + (\mu_* - \mu_i + 2\mathcal{S}_{i,t-1}).
\end{aligned}$$

By definition, $\mathcal{S}_{i,t}$ is the slack quantity $\mathcal{S}_{i,t} = \sqrt{\frac{2\beta \log(t)}{Q_i(t)}}$. By setting $s_i = \frac{8\beta \log(t)}{\Delta_i^2}$, we ensure that the event $Q_i(t-1) > s_i$ implies that

$$Q_i(t-1) > \frac{8\beta \log(t)}{\Delta_i^2} \Leftrightarrow \mu_* - \mu_i + 2\mathcal{S}_{i,t-1} < 0.$$

It follows that

$$\begin{aligned}
&\mathbb{P}[\widehat{\mu}_{i,t-1} - \mathcal{S}_{i,t-1} \leq \widehat{\mu}_{*,t-1} - \mathcal{S}_{*,t-1}, Q_i(t-1) > s_i] \\
&\leq \mathbb{P}[\widehat{\mu}_{*,t-1} - \mathcal{S}_{*,t-1} - \mu_*] \\
&\quad + \mathbb{P}[\mu_i - \widehat{\mu}_{i,t-1} + \mathcal{S}_{i,t-1} - 2\mathcal{S}_{i,t-1}].
\end{aligned}$$

By a union bound,

$$\mathbb{P}[\widehat{\mu}_{*,t-1} - \mathcal{S}_{*,t-1} - \mu_* \geq 0] \leq \sum_{s=1}^{t} t^{-\beta}.$$

Similarly, it also follows that

$$\mathbb{P}[\mu_i - \widehat{\mu}_{i,t-1} + \mathcal{S}_{i,t-1} - 2\mathcal{S}_{i,t-1}] \leq t^{1-\beta}.$$

By assumption, for each $\mathcal{I}_{t,k}$, $\forall i, j \in \mathcal{I}_{t,k}$, it is the case that $i \in N_t^{out}(j)$. Moreover, for any $i \in [K]$, it must be the case that for every t, $i \in \mathcal{I}_{t,k}$ for some $k \in [p]$. Thus, we can write

$$\begin{aligned}
Q_i(t) &= \sum_{s=1}^{t} 1_{i \in N_s^{out}(I_s)} = \sum_{s=1}^{t} \sum_{j=1}^{K} 1_{i \in N_s^{out}(j)} 1_{I_s = j} \\
&\geq \sum_{s=1}^{t} \sum_{j \in \mathcal{I}_{s,k}} 1_{i \in N_s^{out}(j)} 1_{I_s = j} \\
&= \sum_{s=1}^{t} \sum_{j \in \mathcal{I}_{s,k}} 1_{I_s = j},
\end{aligned}$$

which implies that

$$\begin{aligned}
&\sum_{i=1}^{K} \sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1) \leq s_i}] \\
&\leq \sum_{k \in [p]} \left(\max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} \Delta_j \right) \\
&\quad \sum_{t=1}^{T} \sum_{j \in \mathcal{I}_{t,k}} \mathbb{E}\left[1_{I_t=j} 1_{Q_j(t-1) \leq \max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} s_j} \right] \\
&\leq \sum_{k \in [p]} \left(\max_{t \in [T]} \max_{j \in \mathcal{I}_{t,k}} \Delta_j \right) \left(\max_{t \in [T]} \max_{j \in \mathcal{I}_{t,k}} s_j \right).
\end{aligned}$$

Combining the above calculations, applying our definition for s_i, and using the fact that the above analysis holds for any such partition shows that

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[l_t(\xi_{I_t}) - l_t(\xi_*)] \\
&\leq \min_{\substack{p \\ \{\mathcal{I}_{t,k}\}_{t \in [T], k \in p}}} \sum_{k \in [p]} \left(\max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} \Delta_j \right) \left(\max_{\substack{t \in [T] \\ j \in \mathcal{I}_{t,k}}} \frac{8\beta \log(T)}{\Delta_j^2} \right) \\
&\quad + K \frac{\beta}{\beta - 2}.
\end{aligned}$$

□

Theorem 4. *On the set of rounds t where the chosen expert abstains, let $\{\mathcal{I}_{t,j}\}_{t \in [T], j \in [2]}$ be a partition of the expert indices $[K]$ into cliques. Then, the regret of* UCB-D *after T rounds is at most:*

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{I_t}) - L_t(\xi_*)] \leq K \frac{2\beta}{\beta - 2} \\
&+ \max_{i \in E_-} \Delta_i \max_{i \in E_-} \frac{8\beta \log(T)}{\Delta_i^2} + \max_{i \in E_+} \Delta_i \max_{i \in E_+} \frac{8\beta \log(T)}{\Delta_i^2}.
\end{aligned}$$

Proof. We begin by decomposing the regret as follows:

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{I_t}) - L_t(\xi_*)] \\
&= \sum_{i=1}^{K} \sum_{t=1}^{T} \mathbb{E}[L_t(\xi_i) - L_t(\xi_*)] \, \mathbb{E}[1_{I_t=i}] \\
&= \sum_{i=1}^{K} \sum_{t=1}^{T} \mathbb{E}[L_t(\xi_i) - L_t(\xi_*)] \\
&\quad \mathbb{E}[1_{I_t=i} (1_{Q_i(t-1) \leq s_i} + 1_{Q_i(t-1) > s_i})],
\end{aligned}$$

where s_i is a quantity to be determined.

Now, let $\mathcal{L}_{i,t} = \widehat{v}_{i,t} - S_{j,t} - c\mathcal{P}_i^-(t) - S_t$, and notice that

$$\begin{aligned}
&\mathbb{P}[I_t = i, Q_i(t-1) > s_i] \\
&\leq \mathbb{P}[0 \leq \mathcal{L}_{i^*,t} - \mathcal{L}_{i,t}, Q_i(t-1) > s_i] \\
&= \mathbb{P}[0 \leq (\mathcal{L}_{i^*,t-1} - \mu_{i^*}) \\
&\quad + \mu_i - \mathcal{L}_{i,t-1} - 4S_{i,t-1} \\
&\quad + \mu_{i^*} - \mu_i + 4S_{i,t-1}, Q_i(t-1) > s_i].
\end{aligned}$$

Let $s_i = \frac{8\log(T)}{\mu_i - \mu_*}$. Then, $Q_i(t) > s_i$ implies that

$$\begin{aligned}
&\frac{8\log(t)}{\mu_i - \mu_*} < Q_j(t-1) \\
&\Leftrightarrow \frac{2}{\sqrt{Q_j(t-1)}} < \frac{\mu_i - \mu_*}{\sqrt{2\beta\log(t)}} \\
&\Rightarrow \frac{1}{\sqrt{Q_j(t-1)}} + \frac{1}{t} < \frac{\mu_i - \mu_*}{2}\sqrt{\frac{2}{\beta\log(t)}} \\
&\Leftrightarrow S_{i,t-1} + S_{t-1} < \frac{\mu_i - \mu_*}{2}.
\end{aligned}$$

Thus, it follows that

$$\begin{aligned}
&\mathbb{P}[I_t = i, Q_i(t-1) > s_i] \\
&\leq \mathbb{P}[0 \leq (\mathcal{L}_{i^*,t-1} - \mu_{i^*})] \\
&\quad + \mathbb{P}[0 \leq \mu_i - \mathcal{L}_{i,t-1} - 4S_{i,t-1}].
\end{aligned}$$

By the union bound, we can write

$$\begin{aligned}
&\mathbb{P}[0 \leq (\mathcal{L}_{i^*,t-1} - \mu_{i^*})] \\
&= \mathbb{P}[0 \leq \widehat{v}_{i,t} - S_{j,t} + c\mathcal{P}_i^-(t) - S_t - \mu_{i^*}] \\
&\leq \mathbb{P}[0 \leq \widehat{v}_{i,t} - S_{j,t} - \mathbb{E}[1_{yh_{i^*}(x)<0} 1_{r_{i^*}(x)>0}]] \\
&\quad + \mathbb{P}[0 \leq c\mathcal{P}_i^-(t) - S_t - \mu_{i^*} - \mathbb{E}[1_{r_{i^*}(x)>0}]] \\
&\leq \left(\sum_{s=1}^{t} t^{-\beta}\right) + t^{-\beta}.
\end{aligned}$$

Similarly, we can also show that

$$\begin{aligned}
&\mathbb{P}[0 \leq \mu_i - \mathcal{L}_{i,t-1} - 4S_{i,t-1}] \\
&= \mathbb{P}[0 \leq \mu_i - \widehat{v}_{i,t} + S_{j,t} - c\mathcal{P}_i^-(t) + S_t - 4S_{i,t-1}] \\
&\leq \mathbb{P}[0 \leq \mathbb{E}[1_{yh_i(x)<0} 1_{r_i(x)>0}] - \widehat{v}_{i,t} - S_{j,t} \\
&\quad + \mathbb{P}[r_i(x) \leq 0] - c\mathcal{P}_i^-(t) - S_t] \\
&\leq \mathbb{P}[0 \leq \mathbb{E}[1_{yh_i(x)<0} 1_{r_i(x)>0}] - \widehat{v}_{i,t} - S_{j,t}] \\
&\quad + \mathbb{P}[0 \leq \mathbb{P}[r_i(x) \leq 0] - c\mathcal{P}_i^-(t) - S_t] \\
&\leq t^{1-\beta} + t^{-\beta}.
\end{aligned}$$

Going back to the first term, we can decompose the loss over the times that the chosen expert accepts versus abstains:

$$\begin{aligned}
&\sum_{i=1}^{K}\sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)\leq s_i}] \\
&= \sum_{i=1}^{K}\sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)\leq s_i} 1_{r_i(x_t)>0}] \\
&\quad + \sum_{i=1}^{K}\sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)\leq s_i} 1_{r_i(x_t)\leq 0}].
\end{aligned}$$

Now, notice that

$$\begin{aligned}
Q_i(t-1) &= \sum_{s=1}^{t-1} 1_{r_{I_s}(x_s)>0} + 1_{r_{I_s}(x_s)\leq 0} 1_{r_i(x_s)\leq 0} \\
&\geq \sum_{s=1}^{t-1}\sum_{i=1}^{K} 1_{r_{I_s}(x_s)>0} 1_{I_t=i},
\end{aligned}$$

so that

$$\begin{aligned}
&\sum_{i=1}^{K}\sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)\leq s_i} 1_{r_i(x_t)>0}] \\
&\leq \max_{t\in[T]:1_{r_{I_t}(x_t)>0}} \max_{i:1_{r_i(x_t)>0}} \Delta_i \\
&\qquad \max_{t\in[T]:1_{r_{I_t}(x_t)>0}} \max_{i:1_{r_i(x_t)>0}} s_i.
\end{aligned}$$

Similarly, it also follows that

$$\begin{aligned}
Q_i(t-1) &= \sum_{s=1}^{t-1} 1_{r_{I_s}(x_s)>0} + 1_{r_{I_s}(x_s)\leq 0} 1_{r_i(x_s)\leq 0} \\
&\geq \sum_{s=1}^{t-1}\sum_{i=1}^{K} 1_{r_i(x_s)\leq 0} 1_{I_t=i},
\end{aligned}$$

so that

$$\begin{aligned}
&\sum_{i=1}^{K}\sum_{t=1}^{T} \Delta_i \, \mathbb{E}[1_{I_t=i} 1_{Q_i(t-1)\leq s_i} 1_{r_i(x_t)\leq 0}] \\
&\leq \max_{t\in[T]:1_{r_{I_t}(x_t)\leq 0}} \max_{i:1_{r_i(x_t)\leq 0}} \Delta_i \\
&\qquad \max_{t\in[T]:1_{r_{I_t}(x_t)\leq 0}} \max_{i:1_{r_i(x_t)\leq 0}} s_i.
\end{aligned}$$

Combining the above calculations and substituting in the

expression for s_i, we see that

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[L_t(\xi_{I_t}) - L_t(\xi_*)] \\
&\leq \max_{t\in[T]:1_{r_{I_t}(x_t)\leq 0}} \max_{i:1_{r_i(x_t)\leq 0}} \Delta_i \\
&\quad \max_{t\in[T]:1_{r_{I_t}(x_t)\leq 0}} \max_{i:1_{r_i(x_t)\leq 0}} \frac{8\beta\log(T)}{\Delta_i^2} \\
&\quad + \max_{t\in[T]:1_{r_{I_t}(x_t)>0}} \max_{i:1_{r_i(x_t)>0}} \Delta_i \\
&\quad \max_{t\in[T]:1_{r_{I_t}(x_t)>0}} \max_{i:1_{r_i(x_t)>0}} \frac{8\beta\log(T)}{\Delta_i^2} \\
&\quad + K\frac{2\beta}{\beta-2}.
\end{aligned}$$

□

B.2.2. SLEEPING REGRET

In this section, we prove the sleeping regret guarantees of AUER-N and UCB-H along with supporting lemmas and Theorem 7. We start with the theorem for AUER-N.

Theorem 5. *For any $T > 0$, the sleeping regret of* AUER-N *is upper bounded by*

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[\ell_t(\xi_{I_t})] - \min_{\sigma} \sum_{t=1}^{T} \mathbb{E}[\ell_t(\xi_{\sigma(W_t)})] \\
&\leq O\left(\beta \log T \sum_{j=2}^{K} \frac{1}{\mu_j - \mu_{j-1}} \max_t \frac{T_j(t)}{Q_j(t)}\right).
\end{aligned}$$

Proof. We initially analyze the regret by focusing on each expert j. For $[i] = \{1, 2, \ldots, i\}$, let $M_{i,j}$ for $i < j$ denote the number of times expert j was chosen, where some expert in $[i]$ could have been chosen. Then,

$$\mathbb{E}[M_{i,j}] = \sum_{t=1}^{T} \mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset}],$$

where W_t is the set of awake experts. For any $s \in [T]$, we can split the sum over the rounds according to $T_j(t-1)$ the number of times that j was chosen:

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset}] \\
&= \sum_{t=1}^{T} \mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset} 1_{T_j(t-1)\leq s}] \\
&+ \sum_{t=1}^{T} \mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset} 1_{T_j(t-1)>s}] \\
&\leq s + \sum_{t=s+1}^{T} \mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset} 1_{T_j(t-1)>s}].
\end{aligned}$$

Now, if expert j were chosen, that is $I_t = j$, then it must have had the lowest confidence bound: $\widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \leq \widehat{\mu}_{k,t-1} - \mathcal{S}_{k,t-1}$ for all $\xi_k \in W_t$. Let $\xi_{k^*} = \operatorname{argmin}_{\xi_k\in W_t\cap[i]} \widehat{\mu}_{k,t-1} - \mathcal{S}_{k,t-1}$, then

$$\begin{aligned}
&\mathbb{E}[1_{I_t=j} 1_{W_t\cap[i]\neq\emptyset} 1_{T_j(t-1)>s}] \\
&\leq \mathbb{P}[\forall \xi_k \in W_t, \widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \leq \widehat{\mu}_{k,t-1} - \mathcal{S}_{k,t-1}, \\
&\qquad W_t \cap [i] \neq \emptyset, T_j(t-1) > s] \\
&\leq \mathbb{P}[\widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \leq \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1}, \\
&\qquad W_t \cap [i] \neq \emptyset, T_j(t-1) > s], \qquad (1)
\end{aligned}$$

where the last inequality comes from the fact that the probability of an intersection of events is less than the probability of one event. Then we can write the inequality

$$\begin{aligned}
0 &\leq \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} - \widehat{\mu}_{j,t-1} + \mathcal{S}_{j,t-1} \\
&= \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} - \widehat{\mu}_{j,t-1} + \mathcal{S}_{j,t-1} - \mu_{k^*} + \mu_{k^*} \\
&\quad - \mu_j + \mu_j - \mu_i + \mu_i + \mathcal{S}_{j,t-1} - \mathcal{S}_{j,t-1} \\
&= -\mu_{k^*} + \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} + \mu_j - \widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \\
&\quad + (\mu_{k^*} - \mu_i) - (\mu_j - \mu_i) + 2\mathcal{S}_{j,t-1}
\end{aligned}$$

as the sum of the following three terms:

$$\begin{aligned}
&[1] \quad -\mu_{k^*} + \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} \\
&[2] \quad -(\mu_j - \mu_i) + 2\mathcal{S}_{j,t-1} + (\mu_{k^*} - \mu_i) \\
&[3] \quad \mu_j - \widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1}.
\end{aligned}$$

Note that at least one of the three terms in non-negative.

By choosing

$$s = \left\lceil \frac{2\beta\log T}{(\mu_j - \mu_i)^2} \max_t \frac{T_j(t-1)}{Q_j(t-1)} \right\rceil,$$

then $T_j(t-1) \geq s \geq \frac{2\beta\log t}{(\mu_j-\mu_i)^2} \frac{T_j(t-1)}{Q_j(t-1)}$. This leads to the following set of inequalities:

$$
\begin{aligned}
Q_j(t-1) &\geq \frac{2\beta \log t}{(\mu_j - \mu_i)^2} \\
\implies (\mu_j - \mu_i)^2 &\geq \frac{2\beta \log t}{Q_j(t-1)} \\
\implies (\mu_j - \mu_i) &\geq \sqrt{\frac{2\beta \log t}{Q_j(t-1)}} \\
\implies (\mu_j - \mu_i) - 2\mathcal{S}_{j,t-1} &\geq 0.
\end{aligned}
$$

Notice that by definition of the set $[i]$, $\mu_{k^*} \leq \mu_i$, which implies $-(\mu_{k^*} - \mu_i) \geq 0$. Thus,

$$
(\mu_j - \mu_i) - 2\mathcal{S}_{j,t-1} - (\mu_{k^*} - \mu_i) \geq 0,
$$

which is equivalent to term [2] being non-positive. Thus, at least one of the [1] and [3] terms are non-negative. This implies that we can rewrite the right-hand side of Inequality 1 as

$$
\begin{aligned}
&\mathbb{P}[\widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \leq \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1}, \\
&\qquad W_t \cap [i] \neq \emptyset, T_j(t-1) > s] \\
&\leq \mathbb{P}[-\mu_{k^*} + \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} \geq 0] \\
&\quad + \mathbb{P}[\mu_j - \widehat{\mu}_{j,t-1} - \mathcal{S}_{j,t-1} \geq 0].
\end{aligned}
$$

Now by Hoeffding's inequality and the union bound,

$$
\begin{aligned}
&\mathbb{P}\left[-\mu_{k^*} + \widehat{\mu}_{k^*,t-1} - \mathcal{S}_{k^*,t-1} \geq 0\right] \\
&\leq \mathbb{P}\left[\exists s \in [1,t] : -\mu_{k^*} + \widehat{\mu}_{k^*,s} - \sqrt{\frac{2\beta \log t}{s}} \geq 0\right] \\
&\leq \sum_{s=1}^{t} \frac{1}{t^\beta} = \frac{1}{t^{\beta-1}}.
\end{aligned}
$$

A similar argument holds for the [3] term and thus,

$$
\begin{aligned}
&\mathbb{E}[M_{i,j}] \\
&\leq \frac{2\beta \log T}{(\mu_j - \mu_i)^2} \max_t \frac{T_j(t-1)}{Q_j(t-1)} + 2\sum_{t=s+1}^{T} \frac{1}{t^{\beta-1}} \\
&\leq \frac{2\beta \log T}{(\mu_j - \mu_i)^2} \max_t \frac{T_j(t)}{Q_j(t)} + \frac{\beta}{\beta - 2}.
\end{aligned}
$$

Now that we have bounded the expectation of $M_{i,j}$, we turn to the sleeping regret, which can be written as

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{j=2}^{K}\sum_{i=1}^{j-1}(M_{i,j} - M_{i-1,j})(\nu_j - \nu_i)\right] \\
&= \mathbb{E}\left[\sum_{j=2}^{K}\sum_{i=1}^{j-1} M_{i,j}(\nu_{i+1} - \nu_i)\right]
\end{aligned}
$$

by simply regrouping terms. Then, by using the bound on the expectation of $M_{i,j}$,

$$
\begin{aligned}
&\sum_{j=2}^{K}\sum_{i=1}^{j-1}\mathbb{E}[M_{i,j}](\mu_{i+1} - \mu_i) \\
&\leq 2\beta \log T \sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\mu_{i+1} - \mu_i)}{(\mu_j - \mu_i)^2} \max_t \frac{T_j(t)}{Q_j(t)} + O(1).
\end{aligned}
$$

We then apply Lemma 1 to conclude the statement of the theorem.

□

Theorem 6. *For any $T > 0$, the sleeping regret of* UCB-H *is upper bounded by*

$$
\begin{aligned}
&\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{I_t})] - \min_\sigma \sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{\sigma(W_t)})] \\
&\leq O\left(\beta \log T \sum_{j=2}^{K} \frac{1}{\nu_j - \nu_{j-1}} \max_t \frac{T_j(t)}{O_j(t)}\right)
\end{aligned}
$$

Proof. We initially analyze the regret by focusing on each expert j. For $[i] = \{1, 2, \ldots, i\}$, let $M_{i,j}$ for $i < j$ denote the number of times expert j was chosen, where some expert in $[i]$ could have been chosen. Then,

$$
\mathbb{E}[M_{i,j}] = \sum_{t=1}^{T}\mathbb{E}[1_{I_t=j}1_{W_t\cap[i]\neq\emptyset}],
$$

where W_t is the set of awake experts. Suppose for now that expert ξ_j is not the all-abstain function $\xi_\diamond$. We can split the sum over the rounds according to $T_j(t-1)$ the number of times that j was chosen:

$$
\begin{aligned}
&\sum_{t=1}^{T}\mathbb{E}[1_{I_t=j}1_{W_t\cap[i]\neq\emptyset}] \\
&= \sum_{t=1}^{T}\mathbb{E}[1_{I_t=j}1_{W_t\cap[i]\neq\emptyset}1_{T_j(t-1)\leq s}] \\
&+ \sum_{t=1}^{T}\mathbb{E}[1_{I_t=j}1_{W_t\cap[i]\neq\emptyset}1_{T_j(t-1)>s}] \\
&\leq s + \sum_{t=s+1}^{T}\mathbb{E}[1_{I_t=j}1_{W_t\cap[i]\neq\emptyset}1_{T_j(t-1)>s}].
\end{aligned}
$$

Now, if expert j was chosen, that is $I_t = j$, then it must have had the lowest confidence bound $\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}$ for all $\xi_k \in W_t - \xi_\diamond$ and $\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq c$.

Let $\xi_{k^*} = \text{argmin}_{\xi_k \in W_t \cap [i]} C(k)$ where $C(k) = \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}$ if $\xi_k \in W_t \cap [i] - \{\xi_\diamond\}$ and $C(k) = c$ if $\xi_\diamond \in W_t \cap [i]$.

$$
\begin{aligned}
&E[1_{I_t=j} 1_{W_t \cap [i] \neq \emptyset} 1_{T_j(t-1)>s}] \\
&\leq \mathbb{P}[\forall \xi_k \in W_t - \xi_\diamond, \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}, \\
&\qquad \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq c, W_t \cap [i] \neq \emptyset, T_j(t-1) > s] \\
&\leq \mathbb{P}[\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq C(k^*), \\
&\qquad W_t \cap [i] \neq \emptyset, T_j(t-1) > s], \qquad (2)
\end{aligned}
$$

where the last inequality comes from the fact that the probability of an intersection of events is less than the probability of one event. If $C(k^*) = \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1}$, then we can write the inequality

$$
\begin{aligned}
0 &\leq \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} - \widehat{\nu}_{j,t-1} + \widetilde{\mathcal{S}}_{j,t-1} \\
&= \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} - \widehat{\nu}_{j,t-1} + \widetilde{\mathcal{S}}_{j,t-1} - \nu_{k^*} + \nu_{k^*} \\
&\quad - \nu_j + \nu_j - \nu_i + \nu_i + \widetilde{\mathcal{S}}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \\
&= -\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} + \nu_j - \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \\
&\quad + (\nu_{k^*} - \nu_i) - (\nu_j - \nu_i) + 2\widetilde{\mathcal{S}}_{j,t-1}
\end{aligned}
$$

as the sum of the following three terms:

$$
\begin{aligned}
&[1] \quad -\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} \\
&[2] \quad -(\nu_j - \nu_i) + 2\widetilde{\mathcal{S}}_{j,t-1} + (\nu_{k^*} - \nu_i) \\
&[3] \quad \nu_j - \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1}.
\end{aligned}
$$

Note that at least one of the three terms in non-negative. If instead $C(k^*) = c$, then we can write the inequality

$$
\begin{aligned}
0 &\leq c - \widehat{\nu}_{j,t-1} + \widetilde{\mathcal{S}}_{j,t-1} \\
&= c - \widehat{\nu}_{j,t-1} + \widetilde{\mathcal{S}}_{j,t-1} - \nu_j + \nu_j \\
&\quad - \nu_i + \nu_i + \widetilde{\mathcal{S}}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \\
&= \nu_j - \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \\
&\quad + (c - \nu_i) - (\nu_j - \nu_i) + 2\widetilde{\mathcal{S}}_{j,t-1}
\end{aligned}
$$

as the sum of the following two terms:

$$
\begin{aligned}
&[1'] \quad \nu_j - \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \\
&[2'] \quad -(\nu_j - \nu_i) + 2\widetilde{\mathcal{S}}_{j,t-1} + (c - \nu_i).
\end{aligned}
$$

By choosing

$$
s = \left\lceil \frac{2\beta \log T}{(\nu_j - \nu_i)^2} \max_t \frac{T_j(t-1)}{O_j(t-1)} \right\rceil,
$$

then $T_j(t-1) \geq s \geq \frac{2\beta \log t}{(\nu_j - \nu_i)^2} \frac{T_j(t-1)}{O_j(t-1)}$. This leads to the following set of inequalities:

$$
\begin{aligned}
O_j(t-1) &\geq \frac{2\beta \log t}{(\nu_j - \nu_i)^2} \\
&\implies (\nu_j - \nu_i)^2 \geq \frac{2\beta \log t}{O_j(t-1)} \\
&\implies (\nu_j - \nu_i) \geq \sqrt{\frac{2\beta \log t}{O_j(t-1)}} \\
&\implies (\nu_j - \nu_i) - 2\widetilde{\mathcal{S}}_{j,t-1} \geq 0.
\end{aligned}
$$

Notice that by definition of the set $[i]$, $\nu_{k^*} \leq \nu_i$ as well as $c \leq \nu_i$, which implies $-(\nu_{k^*} - \nu_i) \geq 0$ and $-(c - \nu_i) \geq 0$. Thus

$$
(\nu_j - \nu_i) - 2\widetilde{\mathcal{S}}_{j,t-1} - (\nu_{k^*} - \nu_i) \geq 0
$$

and

$$
(\nu_j - \nu_i) - 2\widetilde{\mathcal{S}}_{j,t-1} - (c - \nu_i) \geq 0,
$$

which is equivalent to term [2] and [2'] being non-positive. Thus, at least one of the [1] and [3] terms are non-negative and also [1'] must be non-negative. This implies that we can rewrite the right-hand side of Inequality 2 as

$$
\begin{aligned}
&\mathbb{P}[\widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \leq C(k^*), \\
&\qquad W_t \cap [i] \neq \emptyset, T_j(t-1) > s] \\
&\leq \mathbb{P}[-\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} \geq 0] \\
&\quad + \mathbb{P}[\nu_j - \widehat{\nu}_{j,t-1} - \widetilde{\mathcal{S}}_{j,t-1} \geq 0].
\end{aligned}
$$

Now by Hoeffding's inequality and the union bound,

$$
\begin{aligned}
&\mathbb{P}\left[-\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} \geq 0 \right] \\
&\leq \mathbb{P}\left[\exists s \in [1,t] : -\nu_{k^*} + \widehat{\nu}_{k^*,s} - \sqrt{\frac{2\beta \log t}{s}} \geq 0 \right] \\
&\leq \sum_{s=1}^{t} \frac{1}{t^\beta} = \frac{1}{t^{\beta-1}}.
\end{aligned}
$$

A similar argument holds for terms [3] and [1'] and thus,

$$
\begin{aligned}
&\mathbb{E}[M_{i,j}] \\
&\leq \frac{2\beta \log T}{(\nu_j - \nu_i)^2} \max_t \frac{T_j(t-1)}{O_j(t-1)} + 2 \sum_{t=s+1}^{T} \frac{1}{t^{\beta-1}} \\
&\leq \frac{2\beta \log T}{(\nu_j - \nu_i)^2} \max_t \frac{T_j(t)}{O_j(t)} + \frac{\beta}{\beta - 2}.
\end{aligned}
$$

We now assume that expert $j = j_\diamond$ is the all-abstain expert. Thus if expert $j_\diamond$ is chosen, then $c \leq \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}$ for all

$\xi_k \in W_t - \xi_\diamond$. Let $\xi_{k^*} = \operatorname{argmin}_{\xi_k \in W_t \cap [i]} \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}$ so that

$$\begin{aligned}
&\mathbb{E}[1_{I_t=j_\diamond} 1_{W_t \cap [i] \neq \emptyset}] \\
&\leq \mathbb{P}[\forall \xi_k \in W_t - \xi_\diamond, c \leq \widehat{\nu}_{k,t-1} - \widetilde{\mathcal{S}}_{k,t-1}, W_t \cap [i] \neq \emptyset] \\
&\leq \mathbb{P}[c \leq \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1}, W_t \cap [i] \neq \emptyset].
\end{aligned}$$

We can rewrite the inequality as

$$\begin{aligned}
0 &\leq \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} - c \\
&= -\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} - c + \nu_{k^*}.
\end{aligned}$$

Since by definition of the set $[i]$ and the fact that $j_\diamond > i$, $c - \nu_{k^*} \geq 0$, then $-\nu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} \geq 0$. Similarly as before, by Hoeffding's and the union bound, $\mathbb{P}[-\mu_{k^*} + \widehat{\nu}_{k^*,t-1} - \widetilde{\mathcal{S}}_{k^*,t-1} \geq 0] \leq \frac{1}{t^{\beta-1}}$. Thus, when expert j is the all-abstain function, for $i < j$,

$$\begin{aligned}
\mathbb{E}[M_{i,j_\diamond}] &= \sum_{t=1}^{T} \mathbb{E}[1_{I_t=j_\diamond} 1_{W_t \cap [i] \neq \emptyset}] \\
&\leq \sum_{t=1}^{T} \frac{1}{t^{\beta-1}} \leq \frac{\beta}{2(\beta-2)}.
\end{aligned}$$

Now that we have bounded the expectation of $M_{i,j}$, we turn to the sleeping regret, which can be written as

$$\begin{aligned}
&\mathbb{E}\left[\sum_{j=2}^{K}\sum_{i=1}^{j-1}(M_{i,j} - M_{i-1,j})(\nu_j - \nu_i)\right] \\
&= \mathbb{E}\left[\sum_{j=2}^{K}\sum_{i=1}^{j-1} M_{i,j}(\nu_{i+1} - \nu_i)\right]
\end{aligned}$$

by simply regrouping terms and using Lemma 3. Then by using the bound on the expectation of $M_{i,j}$,

$$\begin{aligned}
&\sum_{j=2}^{K}\sum_{i=1}^{j-1}\mathbb{E}[M_{i,j}](\nu_{i+1} - \nu_i) \\
&\leq \sum_{j=2}^{K} 2\beta \log T \max_t \frac{T_j(t)}{O_j(t)} \sum_{i=1}^{j-1} \frac{1}{(\nu_j - \nu_i)^2}(\nu_{i+1} - \nu_i) \\
&- 2\beta \log T \max_t \frac{P_{j_\diamond}(t)}{O_{j_\diamond}(t)} \sum_{i=1}^{j_\diamond - 1} \frac{1}{(\nu_{j_\diamond} - \nu_i)^2}(\nu_{i+1} - \nu_i) \\
&+ O(1)
\end{aligned}$$

where we added and subtracted the factor

$$2\beta \log T \max_t \frac{P_{j_\diamond}(t)}{O_{j_\diamond}(t)} \sum_{i=1}^{j_\diamond - 1} \frac{1}{(\nu_{j_\diamond} - \nu_i)^2}(\nu_{i+1} - \nu_i)$$

of the all-abstain expert and used the fact that $E[M_{i,j_\diamond}] = O(1)$. Note that $O_{j_\diamond}$ is the number of times the all-abstain expert was observed and $P_{j_\diamond}$ is the number of times it was chosen. We now focus on the first term $\sum_{j=1}^{K-1} 2\beta \log T \max_t \frac{T_j(t)}{O_j(t)} \sum_{i=1}^{j-1} \frac{1}{(\nu_j - \nu_i)^2}(\nu_{i+1} - \nu_i)$, which we can write it as

$$2\beta \log T \sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\nu_{i+1} - \nu_i)}{(\nu_j - \nu_i)^2} \max_t \frac{T_j(t)}{O_j(t)}.$$

We then apply Lemma 1 to conclude the statement of the theorem. □

Lemma 1. *For $\mu_1 \leq \mu_2 \leq \ldots \leq \mu_K$, then*

$$\sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\mu_{i+1} - \mu_i)}{(\mu_j - \mu_i)^2} F_j \leq 2 \sum_{j=2}^{K} \frac{1}{\mu_j - \mu_{j-1}} F_j$$

for some $F_j > 0$.

Proof. Consider the following simple change of variable

$$\begin{aligned}
&\sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\mu_{i+1} - \mu_i)}{(\mu_j - \mu_i)^2} F_j \\
&= \sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\mu - \mu + \mu_{i+1} - \mu_i)}{(\mu - \mu + \mu_j - \mu_i)^2} F_j \\
&= \sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{(\mu - \mu_i - (\mu - \mu_{i+1}))}{(\mu - \mu_i - (\mu - \mu_j))^2} F_j
\end{aligned}$$

where $\mu > \max_i \mu_i$. Letting $\tilde{\mu}_i = \mu - \mu_i$, then $\tilde{\mu}_1 \geq \tilde{\mu}_2 \geq \ldots \geq \tilde{\mu}_K$ and

$$\begin{aligned}
&\sum_{i=1}^{j-1} \frac{(\mu - \mu_i - (\mu - \mu_{i+1}))}{(\mu - \mu_i - (\mu - \mu_j))^2} F_j \\
&= \sum_{i=1}^{j-1} \frac{(\tilde{\mu}_i - \tilde{\mu}_{i+1})}{(\tilde{\mu}_i - \tilde{\mu}_j)^2} F_j
\end{aligned}$$

so that we can directly apply Lemma 2 □

Lemma 2. *For $\mu_1 \geq \mu_2 \geq \ldots \geq \mu_K$, let $F_j > 0$ for all $j \in [1, K]$ and define $\Delta_{i,j} = \mu_i - \mu_j$, then*

$$\sum_{j=2}^{K}\sum_{i=1}^{j-1} \frac{F_j}{\Delta_{i,j}^2} \Delta_{i,i+1} \leq 2 \sum_{j=2}^{K} F_j \Delta_{j-1,j}^{-1}.$$

Proof. The following proof is adapted from Lemma 3 of Kleinberg et al. (2008) to include the factor F_j. By changing the order of summation,

$$\sum_{j=2}^{K}\sum_{i=1}^{j-1}\frac{F_j}{\Delta_{i,j}^2}\Delta_{i,i+1} = \sum_{i=1}^{K-1}\Delta_{i,i+1}\sum_{j:j>i}\frac{F_j}{\Delta_{i,j}^2}.$$

We then analyze the second sum by fixing an expert $i \in [n]$,

$$\begin{aligned}
&\sum_{j:j>i}\frac{F_j}{\Delta_{i,j}^2}\\
&=\sum_{j=2}^{K}1_{j>i}\frac{F_j}{\Delta_{i,j}^2}\\
&=\int_{x=0}^{\infty}\#\{j : j>i, \Delta_{i,j}^{-2}F_j \geq x\}dx\\
&=\int_{x=0}^{\infty}\#\{j : j>i, \Delta_{i,j}F_j^{-1/2} \leq x^{-1/2}\}dx\\
&=-2\int_{y=\infty}^{0}\#\{j : j>i, \Delta_{i,j}F_j^{-1/2} \leq y\}y^{-3}dy\\
&=2\int_{y=0}^{\infty}\#\{j : j>i, \Delta_{i,j}F_j^{-1/2} \leq y\}y^{-3}dy
\end{aligned}$$

where we changed variables of integration $x^{-1/2} = y$.

Thus combining the above,

$$\begin{aligned}
&\sum_{i=1}^{K-1}\Delta_{i,i+1}\sum_{j:j>i}\frac{F_j}{\Delta_{i,j}^2}\\
&=2\sum_{i=1}^{K-1}\Delta_{i,i+1}\int_{y=0}^{\infty}\#\{j : j>i, \Delta_{i,j}F_j^{-1/2} \leq y\}y^{-3}dy\\
&=2\int_{y=0}^{\infty}y^{-3}\left(\sum_{i=1}^{K-1}\Delta_{i,i+1}\#\{j{:}j>i, \Delta_{i,j}F_j^{-1/2} \leq y\}\right)dy
\end{aligned}$$

where we changed the order of integration and summation. We then expand $\#\{\}$ into a sum of indicator functions

$$\begin{aligned}
&2\int_{y=0}^{\infty}y^{-3}\left(\sum_{i=1}^{K-1}\Delta_{i,i+1}\#\{j{:}j>i, \Delta_{i,j}F_j^{-1/2} \leq y\}\right)dy\\
&=2\int_{y=0}^{\infty}y^{-3}\left(\sum_{i=1}^{K-1}\Delta_{i,i+1}\sum_{j=i+1}^{K}1_{j>i,\Delta_{i,j}F_j^{-1/2}\leq y}\right)dy\\
&=2\int_{y=0}^{\infty}y^{-3}\left(\sum_{j=2}^{K}\sum_{i=1}^{j-1}\Delta_{i,i+1}1_{j>i,\Delta_{i,j}F_j^{-1/2}\leq y}\right)dy
\end{aligned}$$

where the last equation is by changing the order of summation. For an expert j and $y \geq 0$, we define

$$i_y(j) := \operatorname{argmin}\{i : i \leq j, \Delta_{i,j}F_j^{-1/2} \leq y\}.$$

In other words, $i_y(j)$ is the the minimum numbered expert $i \leq j$ such that $\Delta_{i,j}F_j^{-1/2}$ is no more than y. Thus,

$$\begin{aligned}
&\sum_{j=2}^{K}\sum_{i=1}^{j-1}\Delta_{i,i+1}1_{j>i,\Delta_{i,j}F_j^{-1/2}\leq y}\\
&=\sum_{j=2}^{K}\sum_{i=i_y(j)}^{j-1}\Delta_{i,i+1}\\
&=\sum_{j=2}^{K}(\mu_{i_y(j)}-\mu_j).
\end{aligned}$$

By using this fact and changing the order of summation and integration

$$\begin{aligned}
&2\int_{y=0}^{\infty}y^{-3}\left(\sum_{j=2}^{K}\sum_{i=1}^{j-1}\Delta_{i,i+1}1_{j>i,\Delta_{i,j}F_j^{-1/2}\leq y}\right)dy\\
&=2\sum_{j=2}^{K}\int_{y=0}^{\infty}y^{-3}(\mu_{i_y(j)}-\mu_j)\\
&=2\sum_{j=2}^{K}\int_{y=\Delta_{j-1,j}F_j^{-1/2}}^{\infty}y^{-3}(\mu_{i_y(j)}-\mu_j)
\end{aligned}$$

where the last equality holds because for values y less than $\Delta_{j-1,j}F_j^{-1/2}$, $i_y(j) = j$ and the integrand is equal to zero. Notice also that $(\mu_{i_y(j)}-\mu_j)F_j^{-1/2} \leq y$, hence

$$\begin{aligned}
&2\sum_{j=2}^{K}\int_{y=\Delta_{j-1,j}F_j^{-1/2}}^{\infty}y^{-3}(\mu_{i_y(j)}-\mu_j)\\
&=2\sum_{j=2}^{K}\int_{y=\Delta_{j-1,j}F_j^{-1/2}}^{\infty}y^{-3}yF_j^{1/2}dy\\
&=2\sum_{j=2}^{K}F_j^{1/2}\int_{y=\Delta_{j-1,j}F_j^{-1/2}}^{\infty}y^{-2}dy\\
&=2\sum_{j=2}^{K}F_j\Delta_{j-1,j}^{-1}
\end{aligned}$$

which is the bound the lemma. □

Lemma 3. *The sleeping regret can be written in terms of*

the conditional expectations in the following way

$$\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{I_t})] - \min_{\sigma}\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{\sigma(W_t)})] = \mathbb{E}\left[\sum_{j=2}^{K}\sum_{i=1}^{j-1}(M_{i,j}-M_{i-1,j})(\nu_j-\nu_i)\right]$$

where $M_{i,j}$ for $i<j$ is the number of times expert j was chosen, where some expert in $[i]$ could have been chosen.

Proof. Let I_t^* be the index of the best ranked expert that is available at time t.

$$\begin{aligned}&\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{I_t})] - \min_{\sigma}\sum_{t=1}^{T}\mathbb{E}[L_t(\xi_{\sigma(W_t)})]\\ &= \mathbb{E}\left[\sum_{t=1}^{T}\sum_{j=2}^{K}\sum_{i=1}^{j-1}1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right]\\ &= \sum_{t=1}^{T}\sum_{j=2}^{K}\sum_{i=1}^{j-1}\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right].\end{aligned}$$

We now analyze $\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right]$ and assume for now that j and i are not the all-abstain expert. Then,

$$\begin{aligned}&\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_j)\right] - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_i)\right]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_j)|r_j>0\right]\mathbb{P}[r_j>0]\\ &\quad - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_i)|r_i>0\right]\mathbb{P}[r_i>0]\end{aligned}$$

where we used the fact that expert i and expert j are not the all-abstain function and hence they are only chosen if they accept, which implies that $\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_i)|r_i\le 0\right] = 0$ and $\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_j)|r_j\le 0\right] = 0$. Since the decision of I_t is independent of the loss once we know whether an arm accepts or abstains,

$$\begin{aligned}&\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_j)|r_j>0\right]\mathbb{P}[r_j>0]\\ &\quad - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_i)|r_i>0\right]\mathbb{P}[r_i>0]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}|r_j>0\right]\mathbb{E}\left[L_t(\xi_j)|r_j>0\right]\mathbb{P}[r_j>0]\\ &\quad - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}|r_i>0\right]\mathbb{E}\left[L_t(\xi_i)|r_i>0\right]\mathbb{P}[r_i>0]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}\right]\mathbb{E}\left[L_t(\xi_j)|r_j>0\right]\\ &\quad - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}\right]\mathbb{E}\left[L_t(\xi_i)|r_i>0\right]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}\right]\left(\mathbb{E}\left[L_t(\xi_j)|r_j>0\right]\right.\\ &\quad \left.-\mathbb{E}\left[L_t(\xi_i)|r_i>0\right]\right)\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}\right](\nu_j-\nu_i).\end{aligned}$$

Thus,

$$\begin{aligned}&\sum_{t=1}^{T}\sum_{j=2}^{K}\sum_{i=1}^{j-1}\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right]\\ &= \sum_{t=1}^{T}\sum_{j=2}^{K}\sum_{i=1}^{j-1}\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}\right](\nu_j-\nu_i).\end{aligned}$$

It is then easy to conclude the statement in the theorem.

Suppose instead that expert j is the all-abstain, then

$$\begin{aligned}&\mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}[L_t(\xi_j)-L_t(\xi_i)]\right]\\ &= \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_j)|r_j\le 0\right]\mathbb{P}[r_j\le 0]\\ &\quad - \mathbb{E}\left[1_{I_t=j}1_{I_t^*=i}L_t(\xi_i)|r_i>0\right]\mathbb{P}[r_i>0]\end{aligned}$$

and the rest of the proof proceed in a similar way as above. □

Theorem 7. *Assume that $\mathbb{P}(r_j(x)>0) = \mathbb{P}(r_{j'}(x)>0)$ for all r_j, r'_j, then $\sum_{j=2}^{K}\frac{1}{\nu_j-\nu_{j-1}} \le \sum_{j=2}^{K}\frac{\mathbb{P}(r_j(x)>0)}{\mu_j-\mu_{j-1}}$. More generally, assume for all $j\in[K]$ such that $\nu_j\neq c$,*

$$\frac{|\,\mathbb{P}(r_j(x)>0)-\mathbb{P}(r_{j-1}(x)>0)|}{\mathbb{P}(r_{j-1}(x)>0)} \le \frac{\nu_j-\nu_{j-1}}{|\nu_j-c|}.$$

Then, the following inequality holds

$$\sum_{j=2}^{K}\frac{1}{\nu_j-\nu_{j-1}} \le \sum_{j=2}^{K}\frac{1}{\mu_j-\mu_{j-1}},$$

where $\mu_1\le\mu_2\le\ldots\le\mu_K$.

Proof. In the first setting, we assume that $\mathbb{P}[r_j(x)\le 0] = \mathbb{P}[r_{j-1}(x)\le 0]$. Thus,

$$\begin{aligned}\mu_j-\mu_{j-1} &= \mathbb{P}[r_j(x)>0](\nu_j-\nu_{j-1})\\ &\le \nu_j-\nu_{j-1},\end{aligned}$$

so that

$$\sum_{j=2}^{K}\frac{1}{\nu_j-\nu_{j-1}} \le \sum_{j=2}^{K}\frac{\mathbb{P}[r_j(x)>0]}{\mu_j-\mu_{j-1}}.$$

In the more general setting, let

$$\delta_{j,j-1} = (\nu_j-c)(\mathbb{P}[r_j(x)>0]-\mathbb{P}[r_{j-1}(x)>0]).$$

Then,

$$\mu_j-\mu_{j-1} = (\nu_j-\nu_{j-1})\,\mathbb{P}[r_{j-1}(x)>0]+\delta_{j,j-1}.$$

By assumption,

$$\begin{aligned}&(c-\nu_j)(\mathbb{P}[r_j(x)>0]-\mathbb{P}[r_{j-1}(x)>0])\\ &\le (\nu_j-\nu_{j-1})\,\mathbb{P}[r_j(x)>0],\end{aligned}$$

Figure 4: Illustration of two types of experts.

so that

$$\begin{aligned}\mu_j - \mu_{j-1} &= (\nu_j - \nu_{j-1})\,\mathbb{P}[r_{j-1}(x) > 0] + \delta_{j,j-1} \\ &\geq \nu_j - \nu_{j-1} \\ &\geq 0.\end{aligned}$$

Thus, $\{\mu_j\}_{j=1}^K$ and $\{\nu_j\}_{j=1}^K$ share the same ordering.

The above calculation also implies that

$$\begin{aligned}&(\nu_j - c)(\mathbb{P}[r_j(x) > 0] - \mathbb{P}[r_{j-1}(x) > 0]) \\ &\leq \mathbb{P}[r_{j-1}(x) > 0](\nu_j - \nu_{j-1}).\end{aligned}$$

Thus, it follows that

$$\begin{aligned}\frac{1}{\nu_j - \nu_{j-1}} &\leq \frac{1}{(\nu_j - \nu_{j-1})\,\mathbb{P}[r_{j-1}(x) > 0] + \delta_{j,j-1}} \\ &= \frac{1}{\mu_j - \mu_{j-1}}.\end{aligned}$$

The final result follows from the observation that if $\mathbb{P}[r_j(x) > 0] = \mathbb{P}[r_{j'}(x) > 0]$ for all j, j', then $\delta_{j,j-1} = 0$.

□

C. Experiments

In this section, we further describe the datasets and results of the UCB-type algorithms.

We extracted the first principal component of the `CIFAR` dataset (Krizhevsky et al., 2009) for the `horse` and `boat` images, projected the images on this component, and normalized the range of the projections to $[-1, 1]$. We generated six types of 1D-experts that are defined by two thresholds θ_1 and θ_2, see Figure 4 for a depiction of two such experts. In this illustration, the first type of expert on the left abstains on sample points with feature values $x < \theta_1$, while it returns a positive label for $\theta_2 > x$ and a negative label for $\theta_1 \leq x \leq \theta_2$. The second type of expert instead abstains for $\theta_2 > x$ and returns a label in the complement. For all six types of experts, the threshold value θ_1 is in $\{-1, -1 + 2/l, -1 + 4/l, \ldots, 1\}$ for $l = 100/6$ and the second threshold value $\theta_2 = \theta_1 + 2/l$.

The features of the synthetic dataset are drawn from the uniform 2D-distribution over $[-1, 1] \times [-1, 1]$ and the label is determined by the sign of the projection of a point onto the normal of the diagonal hyperplane $y = -x$. The experts are defined by linear functions centered at the origin with slopes in the set $\{\pi/l, \pi 2/l, \ldots, \pi\}$ for $l = 20$ and have labels determined by the projection onto the normal of their planes. Each expert also has a confidence-based abstention region, that is, it abstains on points within a certain distance, $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, from the decision surface.

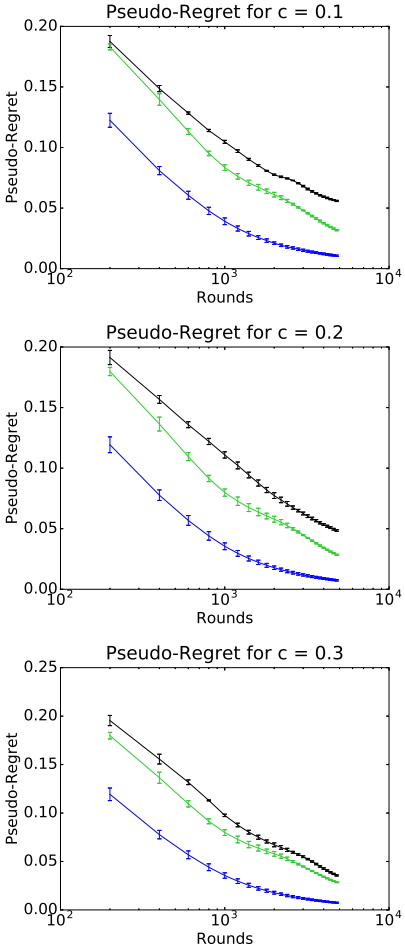

Figure 5: Plot of the pseudo-regret of UCB-type algorithms with 100 experts as a function of the number of rounds for synthetic dataset. UCB-N.ABS is the black line, UCB-D is the green line, and UCB-H is the blue line.

Figure 5 shows the pseudo-regret of the three algorithms on synthetic dataset. The UCB-H algorithm substantially outperforms both the UCB-N.ABS and UCB-D algorithms while the UCB-D algorithm admits a better pseudo-regret than that of the UCB-N.ABS algorithm.